Response to Public Consultation

EDPB Guidelines 02/2026 on Anonymisation

Technical Contribution on Recipient-Bound Release, Cumulative Disclosure and Dynamic Reassessment

Respondent: Sangam Kumar Das
Independent inventor and researcher

Capacity: Individual technical respondent. This submission does not represent any client, employer, institution or commercial party.

Date: 30 July 2026

Executive Summary

I am grateful to the European Data Protection Board for the opportunity to comment on Guidelines 02/2026 on Anonymisation.

I submit this response in an individual technical capacity. It does not offer a legal interpretation of the GDPR. It addresses a narrower engineering question: how an anonymisation assessment can remain operationally effective when data is disclosed repeatedly, automatically and through systems in which the recipient, destination, model, purpose or available auxiliary information may change over time.

The Guidelines provide a valuable framework based on:

1. the perspective of each relevant entity;
2. the means reasonably likely to be used by that entity; and
3. the three criteria of No Record Isolation, No Linkage and No Inference.

The contextual approach is particularly important because the same information may be anonymous for one independent recipient but personal data for another entity possessing different information or technical capabilities.

My contribution concerns the infrastructure needed to preserve the factual assumptions underlying such an assessment.

Where anonymity depends on the capabilities of a particular recipient, it is not enough for the controller merely to document that the data was intended for that recipient. The system should, where appropriate and proportionate, be capable of ensuring that the data is released only to the recipient, destination and processing context against which the assessment was performed.

Similarly, an assessment of a single disclosure may not capture the identifiability created through a sequence of individually low-resolution disclosures. Previous outputs released to

the same recipient may themselves become additional information for the purposes of No Linkage and No Inference.

I therefore invite the Board to consider four clarifications:

1. whether a contextual assessment should consider the controller's ability to bind a disclosure to the recipient and destination actually present at release;
2. whether release-time technical enforcement should be expressly recognised as a measure supporting the continuing reliability of an anonymisation assessment;
3. whether the three criteria should be assessed against cumulative disclosures made to a relevant entity over time; and
4. whether reassessment should be triggered not only periodically, but also by material changes in recipients, models, available datasets, technical capabilities or policy state.

The proposed technical architecture does not assume that access control, encryption or contractual restriction independently makes data anonymous. The legal and technical anonymity assessment must still be performed. The architecture instead seeks to ensure that the result of that assessment is applied at the point where data or a derived output would become externally available.

1. Introductory Remarks

The Guidelines represent a substantial and timely update to the earlier Article 29 Working Party approach. They recognise that identifiability is contextual, that means available through third parties may be relevant, that aggregate information may reveal information about individuals, and that artificial intelligence may reduce the cost and time required for re-identification.

These are technically important developments.

I offer this response in support of the framework. My comments do not challenge the three criteria or the contextual approach. They concern the operational interval between:

- performing an assessment;
- authorising a category of processing or disclosure; and
- the later moment at which a particular disclosure actually occurs.

In automated infrastructure, these events may be separated by time, system layers and organisational boundaries. The assessed recipient may not be the endpoint ultimately selected by a software component. The assessed model may have been replaced or updated. A previously valid purpose or policy may have expired. Earlier outputs may also have increased the information already available to the recipient.

The resulting problem is not necessarily an error in the original assessment. It is a problem of ensuring that the conditions on which the assessment depended remain true when the disclosure becomes effective.

2. Scope of This Response

This response addresses four related matters:

Comment	Subject	Principal provisions
1	Binding a contextual assessment to the actual recipient and destination	Paragraphs 9, 11–13, 20–21, 29–30 and 89–90
2	The interval between assessment and actual release	Paragraphs 29, 34–35, 41–42 and 88–90
3	Cumulative disclosure under No Linkage and No Inference	Paragraphs 17, 55–68 and 79–86
4	Reassessment in rapidly changing AI environments	Paragraphs 35, 84, 88 and 92–93

3. Comment 1 — Binding the Contextual Assessment to the Actual Recipient

Relevant provisions

Paragraphs 9 and 11–13 address access to data and the identification of relevant entities. Paragraphs 20–21 recognise that means of identification may be distributed among different entities. Paragraphs 29–30 require consideration of the release context, access restrictions, receiving entities and unauthorised or malicious actors. Paragraphs 89–90 address accessibility and technical, organisational and contractual restrictions.

The issue

The contextual approach requires the anonymising controller to assess the data from the perspective of the relevant entities for whom it is intended to be anonymous.

This is analytically sound. Its practical application, however, depends on whether the controller can preserve the relationship between:

- the entity whose perspective was assessed; and
- the entity that actually receives or obtains access to the data.

Modern digital systems frequently contain multiple possible release paths, including:

- analytics and advertising software development kits;
- cloud synchronisation or backup functions;
- remote logging systems;
- browser and operating-system services;
- application programming interfaces;
- automated export functions;
- retrieval and tool integrations; and
- autonomous AI agents selecting tools, destinations or external services.

The existence of these systems does not necessarily mean that data will be disclosed unlawfully or to an unexpected party. The technical concern is that an authorisation granted

at one layer may not remain cryptographically or technically bound to the recipient, endpoint and purpose evaluated in the anonymisation assessment.

For example, a controller may correctly assess that a dataset is anonymous from the perspective of an independent research institution because that institution lacks the additional information needed for identification. If an automated workflow instead routes the same information to another service provider possessing cross-platform identifiers, location histories or customer records, the factual basis of the assessment has changed.

The relevant question is therefore not only:

For whom was the information intended to be anonymous?

It is also:

Can the controller ensure, with reliable and enforceable measures, that the particular disclosure is made only to an entity or class of entities covered by the assessment?

The Guidelines already recognise that the context of release and effective restrictions on access are relevant objective factors. They also identify unauthorised and malicious recipients as potentially relevant entities.

The requested clarification would connect these principles more expressly to the technical release path.

Proposed clarification

I respectfully suggest that the final Guidelines include wording substantially equivalent to the following:

Where an assessment under the contextual approach depends on the capabilities, available information or legal independence of a specified recipient or category of recipients, the anonymising controller should consider whether the technical architecture reliably binds the disclosure to that recipient or category. Where the architecture cannot prevent disclosure to an unassessed recipient, endpoint or processing context, the assessment should account for the broader set of entities reasonably likely to obtain access or should use a more conservative approach.

This would not require the Guidelines to mandate any particular technology. It would clarify that the intended recipient and the actual recipient cannot be treated as interchangeable when the identity of the recipient is determinative of the result.

4. Comment 2 — The Interval Between Assessment and Release

Relevant provisions

Paragraph 29 identifies the data's properties, release context, access restrictions, additional information, time, cost and available technology as objective considerations. Paragraph 34 states that relevant measures should be reliable, verifiable and enforceable. Paragraph 35 recommends reassessment as re-identification capability changes. Paragraph 41 requires adequate documentation. Paragraphs 88–90 recognise that entities' capabilities and access may vary over time.

The issue

An anonymisation assessment is normally performed before the processing or disclosure that relies on it.

In many systems, one assessment may govern numerous later operations:

- repeated access to a dataset;
- recurring statistical reports;
- API responses;
- AI-generated outputs;
- location or telemetry disclosures;
- exports to external recipients; and
- automated agent actions.

This creates a potential **time-of-assessment/time-of-release gap**.

Between those two points, one or more relevant conditions may change:

- the intended recipient may have changed corporate control;
- the endpoint may have moved to another jurisdiction or cloud region;
- the approved processing purpose may have expired;
- access may have been revoked;
- a model or analytical method may have changed;
- a new external dataset may have become available;
- the requested resolution may be greater than the resolution assessed;
- a security incident may have occurred; or
- previous disclosures may have altered the recipient's effective capabilities.

A documented assessment remains important, but documentation by itself does not prevent a later operation from proceeding under changed conditions.

The Guidelines correctly state that contractual restrictions should complement technical measures and that relevant measures should be reliable, verifiable and enforceable. They also recognise that technical and organisational measures can restrict access, although such measures do not by themselves make the information anonymous.

A useful further distinction would be between:

1. an assessment whose result is recorded and relied upon through organisational procedures; and
2. an assessment whose essential conditions are technically checked at the moment of disclosure, so that a non-conforming release does not complete.

Release-time enforcement

A release-time system could validate, immediately before disclosure:

- the recipient or requesting entity;
- destination endpoint;
- jurisdiction or processing region;
- authorised purpose;
- data category;
- requested resolution;
- applicable policy version;
- policy expiry;
- revocation status;
- permitted model or processing environment;
- previous disclosures to the recipient; and
- whether the proposed output still passes the applicable anonymity tests.

The system could then:

- release the approved output;
- reduce its resolution;
- suppress particular fields;
- return an aggregate result;
- require a larger cohort;
- introduce an approved privacy transformation; or
- prevent the release.

The measure would not determine, as a matter of law, that the data is anonymous. It would help ensure that a determination already made under the Guidelines is applied to the particular release.

Proposed clarification

I respectfully suggest wording substantially equivalent to:

When evaluating the reliability of technical and organisational measures supporting an anonymisation assessment, controllers should consider whether the relevant conditions are merely documented in advance or are also verified and enforced when the data is accessed or disclosed. Where technically and proportionately feasible, release-time verification may reduce the possibility that data is disclosed under an expired, revoked or materially changed assessment. Such controls support but do not replace application of the three anonymity criteria.

This would give controllers a practical engineering objective while preserving the distinction between access restriction and anonymisation.

5. Comment 3 — Cumulative Disclosure Under No Linkage and No Inference

Relevant provisions

Paragraph 17 recognises that aggregate data may relate to individuals after additional processing. Paragraphs 55–58 address record isolation and aggregation. Paragraphs 59–65 define linkage between records in different datasets. Paragraphs 66–68 require consideration of specific and meaningful inferences based on the given data and additional information. Paragraphs 79–82 address membership inference, inferred linkage, aggregate data, AI models and prompting. Paragraphs 83–86 address advanced AI techniques, dimensionality, resolution and the availability of additional information.

The Guidelines therefore already contain much of the conceptual basis needed to address cumulative disclosure.

The issue

The Guidelines primarily express the three criteria in relation to "the given data" and to other data or information with which it may be combined.

In many operational systems, however, the recipient does not receive a single fixed dataset. The recipient receives a stream of outputs over time.

Examples include:

- repeated location queries;
- recurring aggregate reports;
- statistical dashboards;
- advertising measurements;
- smart-device telemetry;
- transport or mobility outputs;
- health analytics;
- model responses;
- synthetic-data queries; and
- agent-generated sequences of related requests.

A single low-resolution or aggregated disclosure may not permit a sufficiently accurate and meaningful inference about an individual. A sequence of such disclosures may permit one.

For example, repeated coarse location disclosures can cumulatively reveal:

- a probable residence;
- a probable workplace;
- a recurring journey;

- regular attendance at a medical facility;
- regular attendance at a religious institution;
- relationships inferred from co-location; or
- a predictable daily routine.

Similarly, repeated aggregate queries with slightly different populations, dates or filters may enable a recipient to derive individual information through subtraction or differencing.

The important technical point is:

Reducing the resolution of each individual disclosure does not necessarily bound the resolution obtainable from the complete sequence of disclosures.

The same recipient's prior outputs may therefore constitute "additional information" for purposes of the next No Linkage or No Inference assessment.

The Guidelines already recognise that No Linkage depends on other datasets and that No Inference must consider external information, aggregate information, prompting and AI-assisted techniques. They also recognise that dimensionality and the amount of available additional information affect re-identification effectiveness.

An express statement on cumulative disclosures would make this principle easier to apply to continuously operating systems.

Proposed clarification

I respectfully suggest that the final Guidelines state:

Where data or derived outputs are disclosed repeatedly to the same relevant entity, or to entities reasonably likely to combine their holdings, the assessment should consider the cumulative information available after the disclosures rather than assessing each output exclusively in isolation. Previous outputs may constitute additional information for the purposes of No Linkage and No Inference.

The Guidelines could further clarify that relevant factors may include:

- frequency of disclosure;
- duration of access;
- overlap between populations;
- repeated use of related attributes;
- changes in temporal or geographical windows;
- query differencing;
- recipient collusion or common control;
- model memory;
- saved conversation or agent state; and
- the recipient's ability to combine outputs automatically.

Technical feasibility

Cumulative disclosure can be addressed by maintaining protected state associated with:

- the recipient;
- data category;
- purpose;
- cohort;
- resolution;
- disclosure period;
- previously released attributes;
- query history; and
- remaining exposure or privacy allowance.

Before releasing a new output, the system can assess the proposed output together with the recipient's previous authorised outputs.

Depending on the application, it may then:

- permit the output;
- reduce precision;
- enlarge the aggregation group;
- suppress an attribute;
- modify the time window;
- introduce an approved privacy-preserving transformation;
- delay the query;
- require a separate assessment; or
- deny the release.

This need not require permanent tracking of an identified natural person. The protected state may instead be associated with the requesting entity, query class, purpose, dataset or release authority. The implementation should itself be designed to avoid creating a new reusable identifier or linkage mechanism.

6. Comment 4 — Reassessment in Rapidly Changing AI Environments

Relevant provisions

Paragraph 35 observes that re-identification likelihood typically increases over time and recommends periodic reassessment. Paragraph 84 states that the state of the art is continuously evolving. Paragraph 88 recognises that entities' capabilities may change. Paragraphs 92–93 state that AI, particularly agentic AI, may reduce the time and cost required to deploy re-identification techniques and combine information from multiple sources.

The issue

Periodic reassessment is necessary, but a purely calendar-based interval may not correspond to the rate at which the relevant risk changes.

A material change may occur immediately when:

- a relevant external dataset becomes public;
- a security breach exposes auxiliary information;
- the recipient gains access to a stronger analytical model;
- an AI agent is connected to new retrieval sources or tools;
- a model acquires persistent memory;
- a previously separate entity enters the same corporate group;
- the processing purpose changes;
- the disclosure destination changes;
- a privacy-preserving technique is shown to be weaker than expected; or
- a new automated re-identification method becomes practically available.

Agentic AI is particularly relevant because an agent can search, retrieve, correlate, test and refine hypotheses across multiple sources without requiring each intermediate step to be manually initiated.

The Guidelines correctly observe that AI may reduce the time and cost required to access and combine relevant information. The engineering implication is that reassessment should, where feasible, also be **event-triggered**.

Proposed clarification

I respectfully suggest wording substantially equivalent to:

Periodic reassessment should be supplemented, where appropriate, by reassessment following a material change to the data, recipient, access conditions, available additional information, analytical model, processing environment or relevant re-identification capabilities. Controllers should define events that invalidate or suspend reliance on a previous assessment.

The Guidelines could also recognise the value of expiry and fail-closed controls:

Where an automated disclosure relies on a prior anonymisation assessment, technical measures may be used to prevent or restrict disclosure when the applicable assessment or policy is expired, revoked, unavailable or no longer verifiable.

This would not require continuous proof that no undisclosed re-identification technique exists. It would establish a safer operational response when the system can no longer verify that the conditions governing the release remain valid.

7. A Technical Architecture for Release-Time Application of the Guidelines

The following architecture illustrates how the clarifications proposed above could be implemented. It is offered as a technology-neutral example rather than as a request for the Board to mandate a specific design.

7.1 Candidate data operation

Each proposed disclosure is first represented as a candidate operation containing, where applicable:

- data source and version;
- requesting entity;
- intended recipient;
- endpoint or destination;
- purpose;
- jurisdiction;
- model or analytical method;
- requested attributes;
- aggregation level;
- temporal and geographic resolution;
- relevant query history; and
- applicable assessment and policy version.

7.2 Non-effective state

Submission or computation of the operation does not itself authorise disclosure.

The operation remains non-effective until the applicable conditions have been validated. The processing environment should not possess independent authority to bypass the release decision merely because it has calculated an output.

This implements a separation between:

- computing a result; and
- being authorised to make that result externally available.

7.3 Protected validation domain

A protected validation component evaluates the proposed operation against:

- recipient identity and role;
- endpoint;
- purpose limitation;
- jurisdiction;
- revocation and expiry;
- permitted data resolution;
- relevant external information;
- cumulative disclosure state;
- No Record Isolation;
- No Linkage; and
- No Inference.

The protected component may be implemented through hardware isolation, trusted execution, protected processes, cryptographic controls or another architecture capable of preventing ordinary application logic from altering the decision.

No particular implementation should be presumed sufficient solely because it uses encryption, trusted hardware or cryptography. Its effectiveness must be evaluated in the specific context.

7.4 Adaptive output transformation

Where the original requested output would not satisfy the applicable conditions, the system may generate a safer representation by applying measures such as:

- aggregation;
- generalisation;
- suppression;
- cohort enlargement;
- temporal coarsening;
- geographic coarsening;
- noise addition;
- clipping;
- rate limitation; or
- query-budget enforcement.

A system may therefore permit different outputs from the same underlying data for different independent recipients where the contextual assessment supports those differences.

7.5 Scoped release authority

Following successful validation, the system may issue a narrowly scoped authorisation bound to:

- the specific operation;
- approved output;
- recipient;
- endpoint;
- purpose;
- destination;
- policy version;
- validity period; and
- intended release boundary.

The authorisation should not function as a general reusable bearer credential. Copying it to another entity or endpoint should not enable another release.

7.6 Final release-boundary verification

The first boundary at which the data or output becomes externally usable independently verifies the authorisation.

The boundary checks that:

- the approved output has not been altered;

- the intended recipient and endpoint are present;
- the authority has not expired or been revoked;
- the authority has not already been consumed;
- the release remains within the approved scope; and
- any required cumulative state has been updated.

If validation fails, the output is not released or is replaced with an authorised lower-resolution representation.

7.7 Validation evidence

The system may create a protected validation record establishing:

- what operation was assessed;
- what policy and assessment were applied;
- which recipient and endpoint were validated;
- whether cumulative disclosure was considered;
- which transformation was applied;
- when the decision was made; and
- whether the release completed.

Such evidence may support the documentation contemplated by paragraph 41.

The evidence record should not unnecessarily contain the personal data being protected. It should also avoid introducing a stable identifier capable of facilitating later linkage.

7.8 Reassessment and revocation

The architecture may support:

- assessment expiry;
- policy revocation;
- recipient revocation;
- model-version restrictions;
- reduction of permitted precision;
- updated inference tests;
- changed exposure thresholds; and
- suspension after a material security or capability event.

This allows a controller to stop future disclosures without relying solely on downstream contractual compliance.

7.9 Proportionality, Latency and Compatibility with Existing Systems

The proposed architecture should not be understood as requiring the complete anonymisation assessment to be repeated synchronously for every individual network packet, database read or output operation.

A practical implementation may separate two functions.

First, an assessment function performs the substantive evaluation required by the Guidelines, including the application of No Record Isolation, No Linkage and No Inference, the consideration of available additional information, and the establishment of permitted recipients, purposes, data resolutions, models and disclosure conditions.

Second, a release-time function verifies the limited set of conditions that must remain true when a particular disclosure occurs. Depending on the context, these conditions may include:

- the identity or category of the recipient;
- the destination endpoint;
- the authorised purpose;
- the applicable policy and assessment version;
- expiry or revocation status;
- the permitted data resolution;
- the approved model or processing environment; and
- the cumulative disclosure state associated with the recipient or query class.

The release-time function therefore need not rerun every statistical, inference or re-identification test for every transaction. More computationally intensive testing may be performed when the assessment is established and repeated when a material change occurs. The release boundary verifies that the proposed operation remains within the resulting authorised conditions.

The performance impact of such verification will necessarily depend on the implementation, hardware, policy complexity, transaction volume and location of the enforcement component. This contribution therefore does not propose a universal latency figure. It submits only that local verification, cached policy state and narrowly scoped release predicates can make enforcement compatible with operational systems, subject to appropriate testing in the relevant deployment.

The proposal also does not require the immediate replacement of existing or legacy infrastructure. Release-time controls may be introduced incrementally at an existing point through which data becomes externally available, including an API gateway, database proxy, analytics export service, cloud egress control, file-download service, AI tool connector or other functionally equivalent boundary.

A legacy application may continue to perform its existing computation while the authority to disclose the resulting data or output is separated from that application and placed at the release

boundary. Deployment may therefore begin with higher-risk processing and disclosure paths and expand according to the controller's technical capacity, risk assessment and proportionality obligations.

The relevant principle is not that every controller must adopt a particular hardware or cryptographic design. It is that, where the result of an anonymisation assessment depends on conditions that may change between assessment and disclosure, the controller should consider whether those conditions can be verified and enforced at the point at which the disclosure becomes effective.

8. Suggested Text for Possible Inclusion in the Final Guidelines

For convenience, I respectfully propose the following four additions.

A. Recipient-bound contextual assessment

Where anonymity is assessed from the perspective of a specified recipient or category of recipients, the anonymising controller should consider whether reliable technical and organisational measures bind the disclosure to that recipient or category. Where data may be disclosed to an unassessed recipient or processing context through means reasonably likely to be used, that broader access should be taken into account.

B. Release-time verification

An anonymisation assessment may govern repeated or automated disclosures. Where proportionate, controllers should consider technical measures that verify the continued validity of the relevant conditions when each disclosure occurs. These measures may support the reliability of an assessment but do not replace application of the three anonymity criteria.

C. Cumulative disclosure

No Linkage and No Inference should be assessed in light of previous data or outputs disclosed to the same relevant entity, and to other entities reasonably likely to combine their holdings. A sequence of disclosures may enable linkage or inference even where each disclosure, viewed separately, would not do so.

D. Event-triggered reassessment

In addition to periodic reassessment, controllers should identify material changes that require reassessment, including changes to the recipient, processing purpose, available auxiliary information, analytical model, access conditions, disclosure destination, security state or relevant re-identification capabilities.

9. Limitations of the Proposed Architecture

For clarity, I do not suggest that:

- access restriction alone makes personal data anonymous;
- encryption alone constitutes anonymisation;
- recipient authentication proves that the data satisfies the three criteria;
- a validation receipt proves the correctness of an inadequate underlying test;
- release-time control eliminates the need for reassessment;
- source data becomes anonymous for a controller that retains realistic means of identification; or
- technical measures remove the need for an appropriate legal basis while personal data is being processed.

The proposal is narrower.

Where an assessment concludes that a particular output is anonymous from a particular relevant perspective, release-time enforcement can help preserve the factual conditions on which that conclusion depends.

The Guidelines remain the basis for determining whether the output is anonymous.

10. Technical Disclosure and Transparency

The execution-time control concepts described in this response are further developed in my international patent materials, including PCT/IB2026/053385 and related technical disclosures available through WIPO PATENTSCOPE.

I refer to those materials solely to demonstrate that the proposed release-time structures have been described at an implementation level and are not presented merely as an abstract policy aspiration.

I am the inventor and applicant associated with those materials. This is disclosed for transparency.

This submission:

- is made in my individual capacity;
- does not represent a commercial party;
- does not request procurement, licensing or funding;
- does not request endorsement of any patent application;
- does not ask the Board to adopt proprietary terminology; and
- does not suggest that compliance with any patent claim establishes GDPR compliance.

The Board may describe any relevant technical principle in neutral and technology-independent terms.

11. Closing Remarks

The three criteria and the contextual approach provide a strong framework for assessing anonymisation.

My comments concern the operational application of that framework in systems where:

- disclosures occur repeatedly;
- recipients and destinations may be selected dynamically;
- AI systems can combine information automatically;
- the recipient's available information changes over time; and
- no human necessarily reviews each individual release.

The central technical proposition is simple:

Where anonymity depends on the recipient, purpose, resolution, available information and processing context, those conditions should, where technically and proportionately feasible, be validated at the point at which the data or output becomes available to that recipient.

A documented assessment explains why a disclosure is considered permissible. Release-time enforcement helps ensure that the actual disclosure is the disclosure that was assessed.

Similarly, cumulative state allows the controller to evaluate not only what the recipient is receiving now, but what that recipient will possess after the new output is added to previous outputs.

These mechanisms do not replace the legal and technical analysis required by the Guidelines. They provide infrastructure through which a correct assessment can remain effective in automated and AI-mediated systems.

I am grateful to the Board for its work on these Guidelines and for the opportunity to contribute. I remain available should any technical point in this response require clarification.

Sangam Kumar Das
Independent inventor and researcher