ISTITUTO ITALIANO
PER LA PRIVACY E LA
VALORIZZAZIONE DEI DATI

Comments on the EDPB Guidelines on Web Scraping

Istituto Italiano per la Privacy e la Valorizzazione dei Dati

Guidelines 03/2026 on web scraping in the context of generative AI

Version 1.0 adopted on 7 July 2026

22 October 2026

The Istituto Italiano per la Privacy e la Valorizzazione dei Dati (Italian Institute for Privacy and Data Valorisation) recommends targeted clarifications and amendments that make the conditions for lawful web scraping sufficiently clear to guide actual development decisions. The final Guidelines should enable an organisation to determine the sources from which it may collect data, what evidence supports the proposed use, and which measures must operate throughout collection, training, and deployment. They should also identify when collection or use must stop.[1]

The practical value of the guidance depends on distinguishing activities that present materially different risks. Automated acquisition for immediate filtering, retention of a reusable corpus, model training, and disclosure through a deployed service should each receive the legal assessment appropriate to that operation. Public availability does not establish permission to process; equally, scale and automation should not become substitutes for examining necessity, impact, and effective safeguards.

Our proposals address seven points: responsibility for each operation, a workable assessment of legitimate interests, the significance of access restrictions, effective transparency and rights, minimisation and accuracy, incidental collection of special categories of personal data, and the treatment of models and remedies. Several comments seek practical guidance on principles already recognised in the draft. Where IIP proposes a change to the wording or interpretative scope, the relevant section identifies that recommendation expressly. The proposed wording below is intended for insertion in the final Guidelines and distinguishes IIP's recommendations from the legal requirements and judgments on which they draw.

1 Allocate responsibility for the operations determined by each actor

Operational clarification of paragraphs 3 to 5 and 16 to 20

Paragraphs 16 to 20 already require roles to be assessed by reference to the particular processing and the actors' actual involvement. IIP recommends illustrating that approach across the sequence from source selection to deployment. A developer may obtain a corpus from an independent supplier, commission scraping under instructions, or jointly design a collection project. The relevant question is which party determines the purposes and essential means of the particular processing. A contract label, payment, or the fact that both parties benefit from the project cannot answer that question on its own.

The exclusion of data brokers that merely hold datasets should not leave the developer without practical guidance on acquisition. The final text should explain what information a purchaser needs to assess its own use: provenance and collection periods, the categories of sources and personal data, the asserted basis for collection and disclosure, material filtering steps, and arrangements for corrections, objections, and deletion. Existing supplier documentation, sampling, and technical evidence should be reusable where adequate. Reconstructing every historical collection event would often be impossible; accepting an unsupported assurance that data are public would be inadequate.

[1] European Data Protection Board, Guidelines 03/2026 on web scraping in the context of generative AI, version 1.0 adopted on 7 July 2026, for public consultation. Paragraph and example references in this contribution refer to that version.

Istituto Italiano per la privacy e la Valorizzazione dei Dati
Piazza di San Salvatore in Lauro, 13
00186 – Roma
https://www.istitutoitalianoprivacy.it/
Telefono: +39 06 97842491
Fax: +39 0623328983
Email: info@istitutoprivacy.it

Joint controllership should be attributed by reference to the stages that the parties jointly determine. If two organisations jointly select sources but only one later determines an independent deployment, responsibility for that later operation requires a separate analysis. Conversely, allocating scraping to a contractor does not remove a developer's responsibility for instructions that result in excessive collection. Article 28 should govern processing on behalf of a controller, while Article 26 addresses genuine joint determination.[2]

For an independently sourced corpus, the practical decision should follow the evidence. A credible indication of leaked material, collection from sources subject to access restrictions, or an unresolved deletion order requires investigation and appropriate restrictions on processing the affected data. The supplier's failure to identify provenance may prevent the purchaser from demonstrating lawfulness. These consequences should be stated without creating automatic joint responsibility for all participants in a commercial chain.

Add after paragraph 20

Roles and responsibilities should be identified for each relevant operation, including acquisition, filtering, retention, training, and subsequent use. Joint controllership extends to the operations whose purposes and essential means are jointly determined. An organisation acquiring a pre-existing dataset should obtain and assess evidence sufficient to demonstrate the lawfulness of its own processing, taking account of provenance, collection conditions, material safeguards, and any indications of unlawful origin. The depth of verification should reflect the risks and the reliability of the available evidence. Contractual assurances may support, but cannot replace, that assessment. Material unresolved deficiencies should be addressed before the affected data are used.

2 Substantiate reliance on legitimate interests

Operational clarifications to paragraphs 23 and 46 to 61 and proposed amendments to paragraph 58 and Example 7

The draft already recognises Article 6(1)(f) as an available legal basis where its three cumulative conditions are met. IIP supports that approach and recommends further examples of its application. The difficulty of obtaining consent does not establish that those conditions are met, while a commercial objective does not in itself disqualify an interest. In Koninklijke Nederlandse Lawn Tennisbond, the Court recognised that a lawful commercial interest can qualify, while maintaining the separate requirements of strict necessity and balancing. The examples should illustrate how these requirements apply separately for both commercial developers and organisations pursuing research or public benefits.[3]

General-purpose models may require broad and diverse corpora to develop language capabilities, cover different contexts, and avoid unrepresentative results. The Guidelines should recognise that broad web crawling can be necessary and proportionate for a defined development objective. The controller should explain the capabilities sought, the relevant languages and domains, why personal information contributes to those capabilities, and the uses planned or excluded. Exhaustive prediction of every downstream application is unnecessary. The assessment should identify practical collection limits and safeguards compatible with that objective. A material expansion beyond the assessed purpose or capabilities requires review before the additional processing begins, including the applicable purpose compatibility analysis.

Necessity should be assessed against credible alternatives for the defined objective. Synthetic data, licensed material, a narrower corpus, or earlier anonymisation are relevant alternatives where they

[2] Regulation (EU) 2016/679 (GDPR), Articles 4(7) and (8), 5(2), 24, 26, and 28.

[3] CJEU, 4 October 2024, Koninklijke Nederlandse Lawn Tennisbond, C-621/22, EU:C:2024:858, paragraphs 37 to 50 and 57.

offer an equally effective and less intrusive approach. Their suitability depends on coverage, reliability, and the capabilities being developed; none is a universal substitute. Existing development records and suitable tests should support the comparison, without requiring exhaustive experiments with every conceivable alternative. The necessity of broad collection should be assessed on the basis of that evidence. The developer should not define the objective circularly as training on the very dataset whose necessity is in question, or dismiss alternatives solely because unrestricted scraping is cheaper.

The scale of collection matters through its effects: the number and characteristics of affected people, the detail accumulated about them, duration, combination across contexts, and opportunities for disclosure. A diverse corpus with limited personal content may differ substantially from a smaller corpus that reconstructs an individual's private life. IIP therefore proposes amending the final sentence of paragraph 58 so that the number of sources is assessed together with those factors, rather than treated as a conclusive indicator of increased interference.

Paragraph 61 already requires the likelihood of further consequences to be assessed in light of the technical and organisational measures in place and the circumstances of the case. IIP welcomes this approach and recommends supplementing it with practical examples and criteria for its application. These should explain how controllers can connect an identified risk to the data, capabilities, access conditions, and reasonably foreseeable uses of the model, and what evidence is appropriate to assess both its likelihood and the effectiveness of safeguards. The assessment remains prospective and does not depend on evidence of actual harm. The distinct question of anonymity, including the effect of statutory barriers to identification, is addressed in section 7 of this contribution.

Example 7 would benefit from a fuller explanation of the purpose, necessity assessment, and factors supporting the favourable outcome of its balancing exercise. This would illustrate the cumulative conditions already set out in the Guidelines and the separate need to meet the other GDPR requirements. An opt-out provides an additional safeguard but does not establish that the collection was necessary. Mousse confirms that a subsequent opportunity to object does not determine whether processing was necessary at the outset.[4]

Add after paragraph 50 and align paragraphs 51 and 52

For general-purpose development, the controller should describe the capabilities sought, the intended range of uses, the relevant data requirements, and material exclusions with sufficient precision to support the necessity assessment. Broad and diverse web collection may be necessary and proportionate where supported by the development objective and evidence of the limitations of realistic alternatives. Neither its scale nor its untargeted character determines lawfulness on its own. Exhaustive prediction of every downstream application is not required. The controller should compare realistic, less intrusive means capable of achieving the defined objective with equivalent effectiveness, using existing evidence where adequate. Material changes in purpose, sources, or intended capabilities require reassessment before the additional processing begins.

Amend the final sentence of paragraph 58

The number and diversity of sources should be assessed together with the nature, temporal coverage, and combination of the information collected, the individuals affected, and the effectiveness of safeguards. An increase in the number of sources may increase interference, but the assessment should explain the resulting effects on individuals rather than infer the level of interference from that number alone.

[4] CJEU, 9 January 2025, Mousse, C-394/23, EU:C:2025:2, paragraphs 67 to 70.

Supplement paragraph 61 with application criteria and a worked example

> *In applying the likelihood assessment described above, the controller should explain how each identified risk relates to the data, capabilities, access conditions, and reasonably foreseeable uses of the model, and what evidence supports the effectiveness of the safeguards relied upon. For example, where the relevant risk is extraction of personal data, the assessment should compare the access actually provided through a controlled service with access to downloadable model weights, where relevant. Evidence about interface-level filters should be assessed for the access setting in which they operate. The assessment should explain how its conclusion changes when a safeguard cannot be applied or is shown to be ineffective. This balancing assessment remains prospective and does not require evidence of actual harm.*

3 Give technical restrictions their correct legal significance

Clarifications to paragraphs 11, 37, and 64 and Examples 4 and 5

The guidance would benefit from distinguishing an access barrier from a machine-readable instruction and from the legal right that an instruction may express. Authentication restricts access to content. A robots.txt file communicates crawler rules; the Robots Exclusion Protocol itself states that those rules are not a form of access authorisation. CAPTCHA may be used to constrain automated access, but its presence does not establish a single legal consequence for every processing purpose. Explaining these distinctions would help controllers identify the applicable requirements in each case.[5]

IIP recommends retaining a precautionary default of excluding sources that clearly indicate opposition to the proposed collection. The legal explanation should nevertheless identify whether the restriction concerns authorised access, a binding contractual obligation, a copyright reservation, or evidence relevant to fairness and reasonable expectations. A website operator does not generally exercise each person's GDPR rights merely by operating the site. Where a signal conveys an individual's objection through an authorised mechanism, that individual right must be addressed on its own terms.

Public accessibility and the absence of crawler restrictions do not authorise every reuse of personal information. In line with paragraph 63, a platform notice does not, by itself, establish reasonable expectations for every person mentioned in its users' posts. The position of a person who deliberately publishes their own professional biography differs from that of a child identified in someone else's photograph or a person described in a forum discussion. IIP recommends expanding Examples 4 and 5 to explain the weight given to these factors alongside the platform's technical choices.

The description of copyright status in Example 7 would benefit from clarification. Content that may be mined under an applicable exception does not thereby become unprotected by copyright. Whether a copyright licence or a valid text and data mining exception applies is a separate question from whether there is a legal basis under the GDPR. Likewise, compliance with the AI Act's copyright policy and training content summary obligations does not establish lawfulness under Articles 6 or 9 GDPR. These assessments may draw on shared evidence while retaining their distinct legal requirements.[6]

Replace the final bullet of paragraph 37 and align paragraph 64

> *The controller should identify and respect applicable restrictions on access and use. As a precautionary collection measure, sources that clearly indicate opposition to the proposed scraping should normally be excluded. The assessment should distinguish effective access controls, machine-readable preferences, binding legal restrictions, and signals conveying data*

[5] IETF, RFC 9309, Robots Exclusion Protocol, September 2022, sections 1 and 3. The technical distinction does not determine the effect of a restriction under applicable law.

[6] Directive (EU) 2019/790, Articles 3 and 4, including Article 4(3) on reservation of rights. Regulation (EU) 2024/1689 (AI Act), Article 53(1)(c) and (d), on copyright policy and the summary of training content.

subjects' rights. Any uncertainty about the legal effect or scope of a signal should be resolved and the conclusion documented before collection. Absence of a restriction is neither consent nor sufficient evidence of reasonable expectations. The effect on those expectations must be assessed in the actual publication context and for the persons whose information is processed.

Clarify Example 7

Compliance with applicable copyright requirements, including any licence or text and data mining exception relied upon, should be established separately from the GDPR assessment. Content eligible for an exception may remain protected by copyright. A platform notice or rights holder's permission does not, by itself, establish the data subject's reasonable expectations or consent, or satisfy a condition under Article 9(2).

4 Make transparency and rights effective without additional exposure

Proposed revisions of paragraphs 27 to 34 and 66 and clarifications to Examples 2 and 3

The draft recognises in paragraphs 25 to 28 that Article 14(5)(b) may apply where its conditions are shown to be met. However, the statement in paragraph 27 that this exception cannot be relied upon "routinely" creates unwarranted legal uncertainty and should be removed: under Recital 62 GDPR, the availability of Article 14(5)(b) depends on objective factual conditions—such as the sheer scale of unstructured web corpora and the lack of direct contact details—rather than the frequency with which a controller trains AI models. IIP recommends further guidance on its practical assessment, including outside research and archiving. A dataset-level assessment is sensible because in large-scale web scraping, organisations should not be expected to scan unstructured web pages to extract contact details or build identity dossiers solely to locate people whom the processing otherwise need not identify.[7]

The assessment should account for the practical effectiveness and privacy consequences of individual notice, alongside effort and the consequences of non-notification. An obsolete identifier, an uncertain match, or a message that discloses sensitive inferences to the wrong person shows why attempting individual outreach from scraped web text creates disproportionate effort and privacy risk under Article 14(5)(b). These are reasons to design public notice appropriately, not reasons to dispense with transparency. Examples 2 and 3 should state that resolving the notice question does not establish the lawfulness of collecting from the sources described.

Publicly available information about data sources should help people understand and challenge processing. A searchable inventory can serve that purpose, but publishing every scraped URL and collection date under paragraph 31 could expose names, identifiers, health interests, or previously obscure pages, while exposing proprietary trade secrets and enabling targeted data poisoning or extraction attacks by malicious actors. IIP recommends information at a level that remains useful without creating a new personal-data index or undermining security, aligned with the EU legislator's calibrated training data summary framework under Article 53(1)(d) of the AI Act: meaningful category- or domain-type source descriptions, collection periods, crawler identity, and accessible instructions for exercising rights. More detailed provenance should remain available to supervisory authorities upon request. Commercial confidentiality and trade secrets should be accommodated in a manner that preserves meaningful information about data sources and satisfies the applicable transparency requirements. For the same reasons, paragraph 34's recommendation to make Data Protection Impact Assessments (DPIAs) publicly available—even in redacted or summary form—should be deleted: public DPIA disclosure goes beyond Article 35 GDPR, risks exposing trade secrets and security

[7] GDPR, Article 14(5)(b), recital 62, and Article 11(1).

vulnerabilities, and undermines the confidentiality of regulatory discussions between controllers and their Lead Supervisory Authority (LSA) under Articles 35 and 36 GDPR.

Paragraph 29 should set out the conditions and safeguards of Article 11 more fully. The fact that identification is unnecessary for the processing purpose is distinct from the controller being able to demonstrate that it cannot identify the person. The latter condition, notice where possible, and the opportunity for the person to supply enabling information are central to Article 11(2). Its reference to Articles 15 to 20 does not create an exemption from Article 21, although the handling of an objection remains subject to Article 12(2), including demonstrated inability to identify the requester. A controller should use information already available and proportionate information supplied by the requester; an official identity document should not be the default when a URL and other appropriate evidence are sufficient.[8]

Furthermore, the suggestion in paragraph 66(f) that controllers create centralized pre-collection opt-out lists allowing individuals to object across websites by submitting identifying information should be removed, as it would force controllers to collect additional cross-site identifiers in conflict with Article 11 GDPR and expose individuals to heightened privacy risks; pre-collection exclusions at web scale should instead rely on standard, proven, and decentralized machine-readable webmaster protocols (decentralised, machine-readable mechanisms for expressing crawling restrictions and purpose-specific limitations on the use of online content, including robots.txt directives and other signals supported by the relevant crawler or service provider). An optional pre-collection exclusion mechanism should also be distinguished from the statutory right to object. The latter requires the response prescribed by Article 21, including cessation of processing unless the controller demonstrates compelling legitimate grounds overriding the individual's interests, rights, and freedoms, or the processing is necessary for the establishment, exercise, or defence of legal claims. Minimal suppression records may be necessary to prevent reacquisition after a valid request. These records should be used only for that purpose, protected by appropriate security measures, and retained for a justified period. A rights process that deletes a record only for the next crawl to restore it is ineffective.

Revise paragraph 27, add after paragraph 28 and revise paragraph 31 and paragraph 34

The disproportionate effort exception under Article 14(5)(b) GDPR applies whenever its objective factual conditions are met, without any restriction on "routine" reliance. Relying on Article 14(5)(b) does not require controllers to quantify the exact number of identifiable individuals or extract contact channels across unstructured web corpora. The assessment should consider the effectiveness of notice and any additional privacy risks caused by identifying or contacting individuals. Public information must remain meaningful and accessible. Source inventories under paragraph 31 should avoid unnecessarily exposing personal data, trade secrets, or security vulnerabilities through granular URLs or search facilities; category- or domain-type source descriptions aligned with Article 53(1)(d) of the AI Act should be recognised as sufficient. Under paragraph 34, DPIAs should remain confidential and available to supervisory authorities upon request rather than being published in full or summary form.

Replace paragraph 29 and clarify paragraph 66(f)

Article 11(1) does not require a controller to acquire additional identifying information solely to comply with the GDPR where identification is unnecessary for its processing purposes. Reliance on Article 11(2) requires the controller to demonstrate that it cannot identify the data subject, inform the person accordingly where possible, and accept additional information enabling identification for the exercise of the rights under Articles 15 to 20. Article 11(2) does not exempt the controller from

[8] GDPR, Articles 11(1) and (2), 12(2) and (6), and 21(1). Article 12(2) governs refusal to act where inability to identify is demonstrated in the circumstances referred to in Article 11(2).

Article 21. Requests under that provision remain subject to Article 12(2), including the requirement to demonstrate any inability to identify the data subject. Any voluntary mechanism for exclusion before collection should supplement statutory rights, with proportionate verification and measures to prevent reacquisition where required. Paragraph 66(f) should not expect controllers to maintain centralized individual opt-out registries across websites, but should instead rely on standard, decentralized machine-readable exclusion protocols (decentralised, machine-readable mechanisms for expressing crawling restrictions and purpose-specific limitations on the use of online content, including robots.txt directives and other signals supported by the relevant crawler or service provider).

5 Assess minimisation and accuracy at each stage

Operational clarification of paragraphs 35 to 42 and 66(g) to 66(i)

The draft already recommends minimisation measures before and after collection. IIP supports the recognition in paragraph 36 that minimisation does not inherently prohibit training on large volumes of data. Its application should reflect the purpose and characteristics of each operation, including filtering, corpus preparation, training, and use. Some content can be classified only after limited machine processing. Screening and the selection of data for training should therefore have distinct controls, with reasonable and demonstrably effective measures for prevention, detection, and correction. The incidental presence of residual special category data should be addressed through the framework discussed in section 6. Subsequent deletion does not retrospectively establish the lawfulness of an otherwise unlawful collection.

Practical safeguards should match the source and the transformation. Excluding account pages, comments, contact blocks, or unnecessary metadata before extraction may be more effective than removing a name from a retained document. Format-based filters help with structured identifiers but cannot reliably capture every contextual disclosure. Validation should examine material failure modes, including failures involving indirect references, different languages, or vulnerable groups relevant to the corpus. A low average error rate does not resolve a serious failure concentrated in a particular subset.

Synthetic data and pseudonymisation should be evaluated by their actual effects. Replacing names with placeholders may reduce exposure while combinations of details within a narrative still allow individuals to be identified. Generating synthetic data from inputs containing personal data does not remove the legal requirements applicable to that processing, and the legal status of the output requires its own assessment. The final text should recognise effective anonymisation where demonstrated and avoid treating any technique's name as proof of its result. Restrictions on raw-data access and retention remain necessary until the relevant processing no longer involves personal data.

Accuracy should be interpreted by reference to the processing purpose. A historical text may accurately record what a person said at a particular time without accurately describing that person's current circumstances. Correcting a training corpus by silently rewriting the historical statement could itself destroy relevant context. The controller should preserve necessary temporal and source context or exclude unsuitable material, while preventing the service from presenting obsolete or disputed assertions as current facts. If that purpose does not require an identifiable person's history, removing attribution may be the appropriate measure.[9]

An official register can be reliable without being an appropriate source for every AI project. Conversely, literary expression or discussion need not be factually verified as if every sentence were a current statement about an identifiable person. The proposed examples should explain what accuracy means

[9] GDPR, Article 5(1)(c), (d), and (e), Article 16, and Article 25. Accuracy must be assessed having regard to the purposes of processing.

for the operation and the output. A disclaimer that generated text may be wrong does not by itself satisfy the obligation where the system produces inaccurate personal data.

Add after paragraph 38 and at the end of paragraph 42

The controller should specify which personal data are required at each stage and apply reasonable and effective measures to exclude or remove unnecessary information at the earliest practicable point. Screening and the selection of data for training should be governed by distinct controls, with measures to address residual information throughout development and use. Effectiveness should be validated against material error patterns and affected groups. Incidental special category data require the separate assessment described in section 3.2.4 of the Guidelines. Subsequent filtering does not retrospectively establish the lawfulness of an otherwise unlawful collection. Accuracy measures should reflect the purpose, date, and context of the information and the personal data the service is expected to produce. Source reliability alone establishes neither necessity nor lawful reuse.

6 Apply proportionate safeguards to incidental special category data

Interpretative proposal and operational amendments concerning paragraphs 72 to 77

IIP welcomes the EDPB's recognition in paragraphs 70 to 74 that the CJEU's *GC and Others* (C-136/17) framework—assessing obligations within the controller's "responsibilities, powers and capabilities"—applies to the incidental collection of special categories of data, and recommends developing the approach in paragraphs 72 to 75 through a standard based on demonstrable reasonable and proportionate steps and an assessment of safeguards across collection, model development, and deployment. Where the incidental presence of residual data cannot reasonably be prevented or detected despite effective measures, it should be addressed within that assessment. Complete elimination before training should not become a separate, absolute condition for applying the reasoning in *GC and Others*.[10]

The relevant distinction is between incidental processing despite reasonable preventive measures and the deliberate selection or retention of special category data for their training value. General awareness that filters may leave residual information does not, by itself, establish intentional processing. Conversely, labelling a collection incidental cannot justify knowingly preserving identified sensitive passages because they improve performance. The final Guidelines should explain the operational application of the GC and Others framework in light of Article 9 and Meta Platforms v Bundeskartellamt. Neither a legitimate interest under Article 6 nor safeguards alone create an additional exception under Article 9.[11]

Reasonable steps should be concrete and verifiable: excluding sources whose characteristics make the proposed collection inappropriate, applying suitable ingestion filters, testing their performance across relevant languages and contexts, limiting access and retention, and preventing dissemination through model safeguards. The assessment should explain each material measure's contribution and the effectiveness of their combination. No single filter should be required to eliminate every residual occurrence on its own. A low average error rate is nevertheless insufficient where testing reveals a material failure affecting a particular group or source category. Existing validation evidence should remain usable while the assumptions on which it rests remain valid.

Paragraph 74(ii) already addresses deletion immediately after collection or as soon as the data are identified. However, demanding both complete pre-filtering or "immediate" deletion at ingestion and state-of-the-art output filters at deployment imposes an unworkable double compliance burden: unlike

[10] CJEU, 24 September 2019, GC and Others, C-136/17, EU:C:2019:773, paragraphs 45 to 48.

[11] CJEU, 4 July 2023, Meta Platforms and Others v Bundeskartellamt, C-252/21, EU:C:2023:537, paragraphs 69 to 76.

structured identifiers (such as phone numbers or email addresses), special category data (such as political opinions or health status) is inherently contextual, causing real-time pre-filtering across billions of webpages to generate high false-positive rates, while scraped training text is rapidly tokenised and fragmented across computing clusters, making instant isolation and extraction operationally unworkable. IIP therefore proposes an operational formulation requiring reasonable and proportionate steps to identify and remove incidental special category data where practicable, using effective measures available for the affected dataset or model. Where removal is not practicable, the controller should document the impediment and implement an effective and proportionate remedy, including robust output filtering and guardrails where extraction or retraining would be disproportionate or technically unfeasible. Effective filtering should not be treated solely as an interim measure pending removal or retraining. Technical constraints should inform the choice of remedy; the incidental persistence of information under effective safeguards should be distinguished from deliberate exploitation of identified sensitive information.

Where the controller relies on Article 9(2)(e), the assessment must concern the data that the relevant individual has manifestly made public. Publication by a third party, accessibility through an open interface, or a platform's decision to display content is insufficient by itself. Schrems v Meta confirms the limits of using an individual's public statement to justify processing other sensitive information about that individual. The final Guidelines should also address Article 10 separately: data relating to criminal convictions and offences may be processed only under official authority or where authorised by Union or Member State law with appropriate safeguards. An assessment under Article 9 cannot replace an assessment of compliance with these conditions. Data concerning children deserve particular protection but do not constitute a special category merely because of the person's age.[12]

Ingestion and deployment controls should be assessed together. Output safeguards and input/output guardrails can materially reduce disclosure and constitute a primary, proportionate safeguard within the controller's powers and capabilities under *GC and Others*, particularly where training text has already been fragmented across compute clusters or encoded into model weights. Their effect depends on the actual access provided: protection within a controlled service may not extend to downloadable weights. The final Guidelines should recognise a holistic combination of ingestion-phase and deployment-phase measures whose effectiveness is supported by evidence and that provides the required protection, while requiring restriction or suspension of affected processing where those requirements cannot be met.

Proposed clarification of paragraph 73 and amendments to paragraphs 74 and 75

Subject to the conditions set out in paragraph 73, the application of the approach described above should take account of reasonable, proportionate, effective, and demonstrable measures throughout collection, model development, and use. The assessment should take account of the controller's responsibilities, powers and capabilities and the relevant technical constraints. The controller should assess the combined contribution of source selection, ingestion filtering, access and retention controls, model safeguards, and corrective measures. The incidental presence of residual special category data despite such measures, or general awareness of that possibility, does not by itself establish intentional processing. No individual safeguard should be required to eliminate every residual occurrence on its own.

The controller should take reasonable and proportionate steps to mitigate incidental special category data, having regard to the contextual nature of such data, the rapid fragmentation of training corpora across computing systems, and effective technical measures available. Identified impediments should be documented and addressed through appropriate restrictions and corrective action. A holistic combination of reasonable ingestion-phase measures and deployment-phase

[12] CJEU, 4 October 2024, Schrems, C-446/21, EU:C:2024:834, paragraphs 78 to 83. GDPR, Articles 9 and 10 and recital 38.

output filtering constitutes sufficient mitigation under paragraph 74. Safeguards against dissemination may contribute to the combined protection; their effectiveness and limitations must be demonstrated for the access actually provided. They do not create an additional Article 9 exception or authorise deliberate selection or retention of special category data for training.

Add an example after paragraph 75

A developer excludes inappropriate source categories, validates ingestion filters, and operates tested output safeguards. Later testing identifies a residual sensitive passage that the filters could not reasonably have been expected to detect. The controller removes the passage from retained corpora where practicable, prevents reacquisition, and assesses its effects in the model. Where the information has already been tokenised across training clusters or absorbed into model parameters such that extraction or retraining would be disproportionate or technically unfeasible, the controller implements robust output filtering and guardrails that prevent its reproduction. Together with reasonable ingestion-phase measures, these safeguards constitute an effective and proportionate remedy. The assessment examines the legal conditions and the combined effectiveness of those measures. It does not treat the mere existence of an imperfect filter as determinative of the legal assessment or accept continued retention because the identified passage is useful for training.

7 Connect model safeguards to effective rights and proportionate remedies

Operational clarifications to paragraphs 60, 66, and 74(iii) to 77

The final Guidelines should connect corpus governance with the model that is ultimately used or released. Evidence that a source was filtered does not alone establish the absence of personal data in the model. Conversely, the use of personal data during development does not establish that every resulting model or output constitutes personal data. The assessment should distinguish the corpus, intermediate versions, model, interfaces, and outputs, examining relevant extraction and attribution routes under the actual conditions of access. Opinion 28/2024 provides the starting point for this analysis.[13]

The prospective assessment of a processing activity under Article 6(1)(f) and the determination of whether information is anonymous answer different questions. For the latter, IIP recommends giving effect to a specific statutory barrier to identification: means whose use is prohibited for the relevant entity and identification route should be taken into account only where there is evidence of their actual use in relation to the data concerned. Proof of completed re-identification is unnecessary once that conduct has made identifying information or capabilities available. Lawful or otherwise available identification routes remain subject to assessment without proof of actual use. This proposal does not make already identifiable information anonymous or displace preventive security obligations, supervisory powers, or liability for unlawful conduct.

An effective rights process needs to identify what must change when a valid request or verified defect concerns particular material. Depending on the affected processing and the applicable right, the response may involve removal from retained corpora, prevention of reacquisition, correction of retrieval sources, model unlearning, retraining, input/output filtering, or restriction of functionality. These measures are not interchangeable, and where personal data has already been tokenised across training clusters or absorbed into model parameters such that extraction or retraining would be disproportionate or technically unfeasible, robust output filtering and guardrails that prevent the model from reproducing the data should be recognised as an effective and proportionate remedy. The

[13] EDPB, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, 17 December 2024, paragraphs 34 to 45 and 49 to 58.

guidance should permit a choice, supported by evidence, among methods that deliver the required protection.

Paragraph 74(iv) requires a process to monitor model outputs and respond to disclosures. IIP recommends explaining how that process can operate effectively without implying indefinite, centralised retention of every user prompt and response. Such logging could generate a new repository of sensitive information. Testing, privacy-preserving detection, proportionate sampling, user reports, and focused investigation can support monitoring where their combined effectiveness is demonstrated. The controller must be able to react promptly to credible evidence of a disclosure and escalate to stronger measures where necessary. Monitoring itself requires a defined purpose, lawful basis, minimisation, access controls, and retention limits.

IIP supports the regular verification and assessment required by paragraph 76 throughout development and use of the model. The final Guidelines should clarify that reviews must take place at intervals proportionate to the risks, the processing context, and the pace of relevant technical developments. Additional reviews should be triggered promptly by material changes or credible evidence that safeguards are ineffective, including new source categories, weaknesses in filtering, fine-tuning that alters relevant capabilities, changes in access to model weights, or a demonstrated extraction route. Existing evidence may be reused where the review confirms that it remains reliable and applicable. Monitoring technological developments under paragraph 77 should inform both scheduled reviews and the need for additional assessment. A uniform interval should not be imposed on materially different arrangements. A change affecting one dataset, model version, or recipient should lead to a focused review, with existing evidence remaining reusable for unaffected circumstances.

The consequences of unlawful collection also need a clear connection to the affected processing. Subsequent anonymisation does not cure an earlier infringement or displace supervisory powers. Where a model is demonstrably anonymous and its operation does not process personal data, that operation should receive the treatment described in Opinion 28/2024; processing of new inputs or outputs containing personal data during deployment remains subject to the GDPR. Where identifiable information persists, the controller must examine the consequences for continued use and effective remediation. This distinction preserves both accountability for the initial act and recognition of a legally meaningful technical result.[14]

Clarify monitoring under paragraph 74(iv) and supplement the regular verification required by paragraph 76

> *The controller should maintain an effective and proportionate process for detecting and responding to disclosures of special category personal data, using measures appropriate to the model, its interfaces, and the access granted to others. This does not require indiscriminate retention of user interactions. The controller should document the effectiveness and limitations of monitoring and remediation. Where extraction or retraining would be disproportionate or technically unfeasible, robust output filtering and guardrails that prevent the model from reproducing the data should be recognised as an effective and proportionate remedy. If the applicable requirements cannot be met, the affected processing must be restricted or suspended as necessary.*
>
> *The controller should verify and assess the effectiveness of these measures regularly throughout development and use of the model, at intervals proportionate to the risks, the processing context, and relevant technical developments. Material changes or credible evidence of ineffective safeguards should trigger an additional review without waiting for the next scheduled assessment. Each review should document its findings, the continuing validity of evidence relied upon, and any*

[14] EDPB, Opinion 28/2024, paragraphs 109 to 135, particularly paragraphs 113 to 116 and 133 to 135. The scope of its analysis of the consequences of unlawful development is expressly defined in paragraph 111.

corrective measures required. Reuse of existing evidence should depend on confirmation that it remains reliable and applicable.

These amendments would make the Guidelines usable as a sequence of decisions: define the operation and purpose, establish the applicable conditions, apply proportionate controls throughout development and use, demonstrate their combined effectiveness, and respond when the evidence changes. IIP recommends drawing that evidence from existing development records, supplier assessments, legitimate interests assessments, and data protection impact assessments where required. Adequate existing records and validation should be reusable without duplicate forms or tests. The final text should distinguish statutory obligations from recommended means of compliance and allow equivalent measures where their effectiveness is demonstrated. That approach protects individuals while giving responsible organisations a clear basis on which to proceed.