DIGIDATA POLICY REPORT

Restoring Legal Certainty in EU Data Protection Law

Why the Digital Omnibus Matters for the GDPR and the Digital Acquis

POSITION STATEMENT

Contextual identifiability belongs in binding law, supported by anti-circumvention safeguards, auditable assurance and disciplined supervision.

Dr M.R. Leiser (Mark)
DigiData Consulting, Ltd.
18 August 2026

POSITION AT A GLANCE

01 **ARTICLE 4**
Codify actor-relative identifiability and realistic downstream responsibility.

02 **ARTICLE 9**
Stop inference and incidental sensitivity from becoming an internal super-trigger.

03 **EPRIVACY**
Replace repetitive consent with a tightly bounded low-risk whitelist.

04 **AI PATHWAYS**
Create safeguarded routes for training and residual special-category data.

05 **INSTITUTIONS**
Bind the EDPB and EDPS to cooperation, evidence and reasoned method.

06 **INNOVATION**
Apply a bounded Innovation Mandate inside the lawful choice set.

Introduction

The Digital Omnibus remains the only live legislative file that can stabilise the trigger conditions of EU data protection law in the short term. The issue is not whether the Union should protect fundamental rights. It plainly should. The issue is whether the rules that activate the GDPR remain clear, democratically settled and proportionate enough for those rights to operate effectively across the Union's digital economy.

As of 31 July 2026, the general Digital Omnibus remains an open ordinary legislative procedure rather than current law. The Cyprus Presidency failed to secure a Council mandate after a blocking group rejected a compromise that no longer delivered credible simplification.[1] Ireland did not simply carry that text forward. It reopened the substantive discussion through a questionnaire on pseudonymisation, cookie consent, AI processing and compliance burdens.[2] Parliament has published a joint ITRE and LIBE draft report and a large body of amendments, but has not formed a negotiating position.[3] The file has therefore entered a political reset, not a settlement.

Council officials then used the 16 July working-group meeting to test four negotiating clusters: identifiability and pseudonymisation, AI development, compliance burdens and ePrivacy exemptions. The discussion revealed a stronger reform constituency than the failed Cyprus text suggested, but not a single coalition. Several delegations supported Article 29a as the politically safer vehicle while demanding binding criteria, objective factors or implementing acts; others preferred EDPB guidance. The Presidency requested written comments by 28 July and scheduled the next working party for 11 September. The intervening period is therefore a drafting window in which the form of the settlement may matter as much as its substantive direction.[4]

The legal argument begins with *EDPS v SRB*. The Court confirmed that identifiability is contextual and actor-relative: the same pseudonymised dataset may remain personal data for the actor that holds, or can realistically obtain, the key, yet fall outside the GDPR for a constrained recipient that lacks means reasonably likely to identify the natural person.[5] The EDPB has now accepted that core proposition in Guidelines 02/2026. It states that anonymity may vary between entities and that the contextual approach reflects the full legal standard.[6] The doctrinal dispute has therefore moved. The live question is no longer whether contextuality exists, but how legislation should bound, evidence and safeguard it.

That threshold now performs work far beyond the GDPR. It determines whether the AI Act, the Data Act, the European Health Data Space, common European data spaces and AI development operate alone or beneath a cumulative GDPR overlay. It also determines how ePrivacy's independent terminal-equipment rule interacts with downstream personal-data processing. An unstable Article 4 trigger therefore generates repeated classification costs across the wider digital *acquis*.

The familiar opposition between simplification and rights protection conceals the decisive institutional choice. When legislation leaves trigger conditions under-specified, guidance, coordinated opinions, enforcement practice and litigation determine the practical perimeter. That arrangement gives formally non-binding instruments much of the market-shaping force of legislation without equivalent democratic settlement, reliance protection or direct review.

The status quo is not neutral. Article 4 leaves organisations to infer the consequences of contextual identifiability from case law and supervisory guidance. Article 9 risks becoming an internal super-trigger whenever sensitivity

[1] Mikołaj Barczentewicz, 'GDPR reform alive again? July update' (How EU Law Influences Tech, 30 July 2026); European Digital Rights, 'The Digital Omnibus is going on summer break. Your rights are not' (16 July 2026); Joint Business Statement, 'Stop the Council Compromise on Digital Omnibus!' (June 2026).

[2] Agence Europe, 'Digital Omnibus – Irish Presidency consults EU countries on pseudonymisation, data processing for AI and cookies' (15 July 2026); Barczentewicz, 'GDPR reform alive again? July update'.

[3] European Parliament, Legislative Observatory, Procedure 2025/0360(COD), 'Simplification of the digital legislative framework – Digital Omnibus (Omnibus VII)', status accessed 31 July 2026; European Parliament, Draft Report PE786.818v01–00, 22 June 2026; European Parliament, IMCO Amendments PE791.061v01–00, 16 July 2026.

[4] Council working-group discussion summary, 16 July 2026 (on file with the author). The Presidency requested written comments by 28 July 2026 and scheduled the next AGS meeting for the morning of 11 September 2026.

[5] *Case C-413/23 P European Data Protection Supervisor v Single Resolution Board EU*:C:2025:645, especially paras 82–100.

[6] European Data Protection Board, Guidelines 02/2026 on Anonymisation, version 1.0, adopted for public consultation on 7 July 2026, especially the Executive Summary and sections 2.2, 3 and 4.

can be inferred or appears incidentally. Article 22 remains awkwardly aligned with lifecycle-based AI Act governance. Scientific research remains exposed to national divergence. Article 5(3) ePrivacy continues to treat consent as the default response to an exceptionally broad technical trigger. AI training remains dependent on general clauses and case-specific supervisory positions.

The Commission correctly recognised that these questions require legislative attention, but its proposal did not go far enough. Its Article 4 drafting needed stricter anti-circumvention rules; Article 41a placed too much scope-affecting power in implementing acts; its cookie reforms created an unstable split between personal and non-personal terminal information; and it left the supervisory architecture largely untouched. The Cyprus direction then moved too far in the opposite direction by deleting legislative answers and returning practical control to EDPB soft law. The Irish reset creates an opportunity to construct a third route: legislative contextuality, operational safeguards, usable assurance and institutional discipline.

Now that the EDPB has accepted the core contextual principle, the remaining task is to put that principle into law with stronger safeguards and greater legal certainty than either the Commission proposal or the EDPB Guidelines presently provide.

A credible settlement should codify actor-relative identifiability in Article 4; preserve responsibility for designed disclosure, outsourced identification and realistic onward transfer; create a statutory assurance framework linked to evidence, codes and certification; confine Article 9 AI derogations to incidental and residual sensitive data with lifecycle safeguards; align Article 22 with meaningful human intervention and AI Act governance; harmonise the research pathway; replace repetitive cookie consent with a purpose-based low-risk whitelist and reusable choice signals; and subject major supervisory positions to stronger duties of cooperation.[7] Within the remaining lawful choice set, an Innovation Mandate should require evidence, alternatives, proportionality and public reasons.[8]

The GDPR now acts as the hinge of the EU digital *acquis*. Article 4 does not merely define a term. It decides when a wider regulatory architecture engages, which actors bear cumulative duties and whether privacy-preserving design produces a legally reliable change of status. That makes the definition of personal data a rule for AI development, data sharing, research, health-data reuse, terminal-equipment access and cross-border infrastructures.

Section 1 examines the political window after the Council reset and Parliament's amendment phase. Section 2 restores the GDPR's risk-based logic. Section 3 addresses institutional balance and integrates a bounded Innovation Mandate and a duty to cooperate. Section 4 reads *EDPS v SRB* together with Guidelines 02/2026 and tests the resulting trigger against real architectures. Section 5 recalibrates Article 9 and its inference-based boundary. Section 6 confronts the remaining pressure points, including ePrivacy, research, automated decisions and AI training. Section 7 explains the cross-*acquis* consequences. Section 8 sets out a legislative blueprint that treats the Omnibus as an opening for reform rather than a sufficient settlement.

Legislative Position as of 31 July 2026

The legal baseline remains the GDPR, the ePrivacy Directive and the AI Act as amended by Regulation (EU) 2026/1744.[9] The general Digital Omnibus, procedure 2025/0360(COD), has not reached a Council general

[7] M R Leiser, 'Recalibrating the EDPB and the EDPS Under a Duty to Cooperate: An Institutional Case for Cooperative Digital Administration' (policy paper, July 2026) <https://osf.io/preprints/lawarchive/4mxz6_v1> accessed 31 July 2026.

[8] M R Leiser, 'The Innovation Mandate: Making Europe's Digital Supervisory State Accountable for Rights, Growth and Scale' (working paper, July 2026) <https://osf.io/preprints/socarxiv/bmcnv_v1> accessed 31 July 2026.

[9] Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744, 24 July 2026. It entered into force on 27 July 2026.

approach, a Parliament negotiating mandate or interinstitutional negotiations.[10] Presidency texts, committee drafts and amendments reveal bargaining positions; they do not create present obligations.

Instrument or file	Position on 31 July 2026	Practical consequence
Current law	The GDPR and ePrivacy Directive remain the baseline; the AI Act now includes Regulation (EU) 2026/1744.	General Omnibus amendments create no obligations until adopted and in force.
General Digital Omnibus 2025/0360(COD)	Ordinary legislative procedure remains open; no Council general approach and no Parliament mandate.	Core GDPR and ePrivacy provisions remain politically contestable.
Council under the Irish Presidency	Cyprus failed to secure a mandate. Ireland reopened the file, tested four negotiating clusters on 16 July and scheduled the next AGS meeting for 11 September.	The Council now has a defined drafting window and an issue-specific, rather than unitary, reform coalition.
European Parliament	Joint ITRE/LIBE draft report, committee opinions, amendments and a targeted impact assessment remain under negotiation.	The draft report is a holding position, not Parliament's final view.
EDPB Guidelines 02/2026	Adopted for consultation on 7 July; consultation closes 30 October 2026.	Contextual anonymity is now the Board's stated legal standard, but the Guidelines create no binding reliance mechanism.
Digital Omnibus on AI	Regulation (EU) 2026/1744 was published on 24 July and entered into force on the third day following publication.	The AI file has moved to implementation while the data-law file remains open.

The separate Digital Omnibus on AI has completed the legislative process. Regulation (EU) 2026/1744 was published in the Official Journal on 24 July 2026 and entered into force on the third day following publication. The general data-law file now moves against that changed baseline: one Omnibus has entered implementation while the GDPR and ePrivacy reforms remain politically open.[11]

Method and Evidential Base

The analysis draws on the Commission proposal and Staff Working Document; the Cyprus Presidency compromise texts; the German-Swedish Article 29a proposal; the Irish Presidency's reported questionnaire and policy programme; the 16 July Council working-group discussion; the joint ITRE and LIBE draft report; IMCO amendments; the EDPB and EDPS Joint Opinion 2/2026; EDPB Guidelines 02/2026 on Anonymisation and the consultation materials surrounding them.[12] It also uses the Court's judgments in *EDPS v SRB* and *WhatsApp Ireland v EDPB*; five documented privacy-preserving architectures; and the Innovation Mandate.[13] and the proposed reform of the EDPB and EDPS under a duty to cooperate.[14] Negotiating documents and public statements are treated as evidence of institutional positions, not as law.

[10] European Parliament, Legislative Observatory, Procedure 2025/0360(COD), 'Simplification of the digital legislative framework – Digital Omnibus (Omnibus VII)', status accessed 31 July 2026; European Parliament, Draft Report PE786.818v01–00, 22 June 2026; European Parliament, IMCO Amendments PE791.061v01–00, 16 July 2026.

[11] Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744, 24 July 2026. It entered into force on 27 July 2026.

[12] EDPB, Public Consultation on Guidelines 02/2026 on Anonymisation, consultation period 8 July to 30 October 2026.

[13] M R Leiser, 'The Innovation Mandate: Making Europe's Digital Supervisory State Accountable for Rights, Growth and Scale' (working paper, July 2026) <https://osf.io/preprints/socarxiv/bmcnv_v1> accessed 31 July 2026.

[14] M R Leiser, 'Recalibrating the EDPB and the EDPS Under a Duty to Cooperate: An Institutional Case for Cooperative Digital Administration' (policy paper, July 2026) <https://osf.io/preprints/lawarchive/4mxz6_v1> accessed 31 July 2026.

Source Note

Public sources were checked against the versions available on 19 August 2026. Regulation (EU) 2026/1744 and EDPB Guidelines 02/2026 are cited according to their current status. Council working documents, meeting summaries, and internal analyses unavailable through a public repository remain identified as "on file with the author".

Section 1: Political Context — A Reopened but Unsettled Legislative Window

Europe's digital regulatory landscape has expanded rapidly over the past decade, producing a dense and increasingly interdependent legal framework. The General Data Protection Regulation[15] now operates alongside the Data Governance Act, the Data Act[16], the Digital Services Act[17], the Digital Markets Act[18], the Artificial Intelligence Act[19], and an emerging architecture of common and sectoral European data spaces, including the European Health Data Space. These instruments do not operate as self-contained silos. They sit within a broader strategic trajectory first articulated in the European Strategy for Data and more recently updated through the Data Union Strategy, both of which seek to combine rights protection, data availability, industrial policy, and digital sovereignty.[20] The result is not simply more regulation. It is a more layered form of regulation in which uncertainty in one foundational regime can transmit friction across the wider digital *acquis*.

Within that architecture, the GDPR retains a structurally prior role. It is not merely another horizontal instrument among many. In practice, it often determines whether actors may use, share, repurpose, combine, or transfer data at all. That matters because the Union's newer digital instruments, whether directed at data intermediation, business-to-business access, platform governance, AI systems, or sectoral re-use frameworks, presuppose some sufficiently stable account of when the GDPR is triggered and what follows once it is.[21] Where that trigger becomes unstable, the problem does not remain confined to data protection compliance. It migrates into AI development, research infrastructures, digital health, industrial data use, and cross-border business models. In that sense, the GDPR functions as a legal hinge of the wider digital framework, and uncertainty around its scope produces system-level rather than merely sector-specific effects.

Figure 1.1. The GDPR as the Hinge of the EU Digital *Acquis*

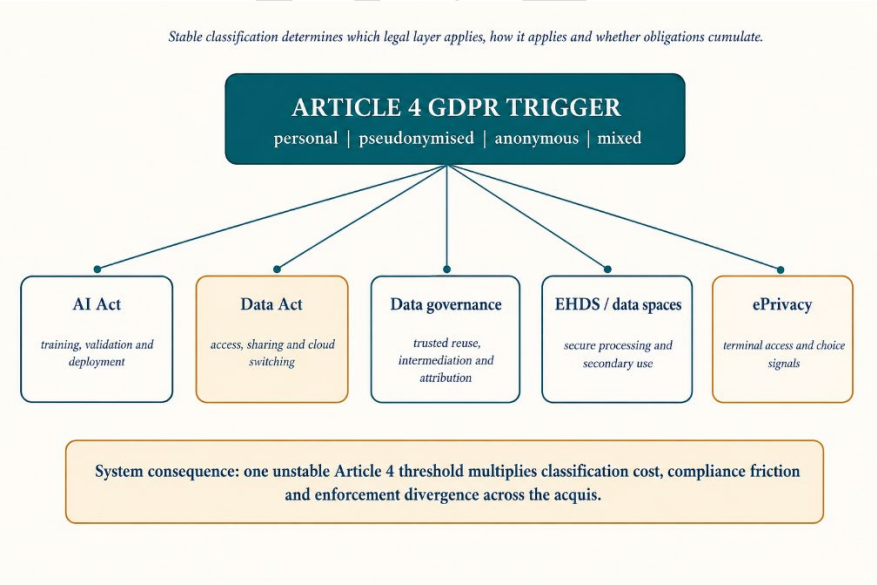

This changed environment is politically important because the EU adopted the GDPR before the present phase of digital regulation had fully taken shape. In 2016, lawmakers conceived the Regulation as a central but

[15] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) [2016] OJ L 119/1.
[16] Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act) [2022] OJ L 152/1.
[17] Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data (Data Act) [2023] OJ L 2023/2854.
[18] Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services (Digital Services Act) [2022] OJ L 277/1.
[19] Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector (Digital Markets Act) [2022] OJ L 265/1.
[20] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 2024/1689; European Commission, A European Strategy for Data, COM(2020) 66 final; European Commission, Data Union Strategy – Unlocking Data for AI, COM(2025) 835 final.
[21] GDPR arts 2(1), 4(1), 4(2), 5 and 6; European Commission, Guidance on the Regulation on a framework for the free flow of non-personal data in the European Union, COM(2019) 250 final, 29 May 2019, 8–10.

relatively self-standing instrument of data protection law. It now operates inside a far wider and more ambitious regulatory ecosystem. The practical significance of questions such as identifiability, anonymisation, legitimate interest, scientific research, automated decision-making, and the treatment of inferred or indirectly revealing data is therefore much greater than it was at the point of adoption. Those questions no longer affect only privacy compliance. They affect the operational viability of a much broader digital rulebook and, with it, the effectiveness of the Union's competitiveness agenda.

That structural tension is now visible in the current legislative cycle. The Digital Omnibus proposal, presented by the Commission on 19 November 2025, reopens a large corpus of the EU digital *acquis* and includes targeted amendments to the GDPR, the ePrivacy Directive, the Data Act, the Data Governance Act and other instruments.[22] The accompanying Staff Working Document supplies the Commission's evidence base.[23] The Commission states that the package should improve legal clarity, reduce administrative burden and align the digital framework with wider competitiveness and implementation objectives.[24] Its 2026 Work Programme places that exercise inside a broader political effort to make the EU digital rulebook more operational and coherent.[25] The EDPB and EDPS contest several elements of that programme, which confirms that the file also reallocates interpretative authority.[26]

Figure 1.2. A Reopened but Unsettled Legislative Window

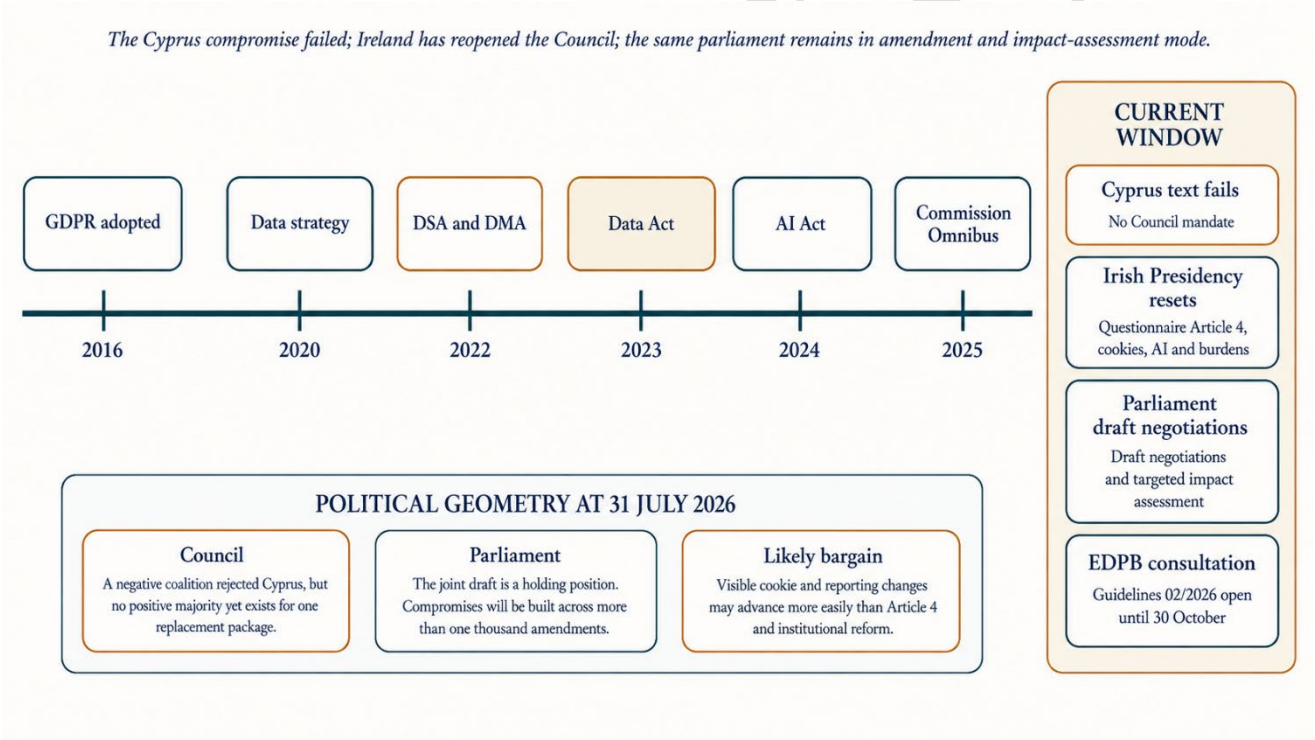

By 31 July 2026, the Council file no longer followed the linear path assumed in the June version of the analysis. Cyprus could not secure agreement on a compromise that deleted the Commission's Article 4 clarification, removed Article 41a and Article 88c, narrowed the AI provisions and left much of the operational work to the EDPB. A blocking group, reportedly including Germany, Italy, Poland, Sweden, Denmark and others, prevented

[22] European Commission, 'Proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2016/679, (EU) 2018/1724, (EU) 2018/1725, (EU) 2023/2854, (EU) 2024/1689 and Directives 2002/58/EC, (EU) 2022/2555 and (EU) 2022/2557 as regards the simplification of the digital legislative framework, and repealing Regulations (EU) 2018/1807, (EU) 2019/1150, (EU) 2022/868, and Directive (EU) 2019/1024 (Digital Omnibus)' COM(2025) 837 final.

[23] European Commission, 'Commission Staff Working Document accompanying the documents Proposal for a Regulation … as regards the simplification of the digital legislative framework (Digital Omnibus) and Proposal for a Regulation … as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)' SWD(2025) 836 final.

[24] European Commission, 'A Simpler and Faster Europe: Communication on Implementation and Simplification' COM(2025) 47 final.

[25] European Commission, 'Commission Work Programme 2026: Europe's Independence Moment' COM(2025) 870 final.

[26] European Data Protection Board and European Data Protection Supervisor, 'Joint Opinion 2/2026 on the Proposal for a Regulation as Regards the Simplification of the Digital Legislative Framework (Digital Omnibus)' (11 February 2026).

the text from advancing. That event matters procedurally because the failed compromise cannot be treated as a *de facto* Council position.[27]

The blocking group nevertheless formed a negative coalition rather than a settled reform majority. Member States could agree that the Cyprus text delivered too little simplification while disagreeing over the remedy. Some may support an operative actor-relative Article 4 test; others may prefer recital clarification, an assurance mechanism or no amendment to the definition. The same fragmentation appears across AI training, special-category data, scientific research and cookies. Opposition to Cyprus therefore created bargaining space without supplying a ready-made replacement package.

The Irish Presidency used that space to reopen the file. Its questionnaire reportedly asks governments to revisit pseudonymisation, data processing for AI, cookie consent and broader compliance burdens.[28] That choice amounts to a procedural reset because it seeks fresh political instructions rather than merely editing the last Cyprus text. Ireland's wider programme combines competitiveness and simplification with an express commitment to fundamental rights.[29] That combination gives it a plausible basis for brokering a guarded legal-certainty settlement rather than a binary deregulatory bargain.

The 16 July Council discussion turns that reset into a more precise political map. Article 4(1) remains the principal red line, especially for delegations that regard the definition of personal data as a constitutional pillar or consider the Omnibus too narrow a vehicle for changing it. Yet the same discussion showed broad acceptance of the underlying contextual problem. Poland, Latvia, Malta, Slovenia, Greece and Cyprus preferred an Article 29a route; Germany and Sweden defended relative anonymity through Article 29a; Denmark supported their initiative; and Italy, France and Spain remained open to further operative clarification. The politically available compromise therefore lies less in whether contextuality exists than in where the legislature places it and what legal effect it gives the rule.[30]

The second axis concerns institutional control. Czechia, Poland and Hungary questioned whether EDPB guidance or an Article 64 opinion could deliver harmonisation. At the same time, the Commission again preferred implementing acts because comitology and judicial scrutiny give them clearer legal consequences. France asked the Commission to explain that choice in writing. Estonia, Cyprus and several more cautious delegations remained more receptive to EDPB-led clarification. The Council is therefore not choosing only between Article 4 and Article 29a. It is choosing among legislation, Commission implementation and supervisory soft law. That is precisely the institutional question that the Omnibus must settle.

The same meeting exposed issue-specific coalitions that Ireland can use to assemble a package. Germany, Poland, Denmark, Austria, Czechia and the Commission favoured restoring an operative AI provision or a comparably binding route; Belgium, the Netherlands, Malta and Cyprus preferred recital 33a, with Finland and Spain showing greater flexibility. On ePrivacy, Estonia, Poland, France, Czechia, Denmark and Italy supported wider low-risk exemptions or serious consideration of contextual advertising, fraud prevention and audience measurement; Spain, the Netherlands, Austria and Hungary pressed for narrower drafting, safeguards or an impact assessment. These are negotiable differences, not a binary divide between rights and simplification.[31]

The next Presidency compromise will therefore carry unusual weight. It will show whether Ireland can turn a coalition against the Cyprus draft into a coalition for legislative contextuality, usable ePrivacy exemptions and proportionate AI pathways. It will also reveal whether the Council intends to settle foundational questions in

[27] Mikołaj Barczentewicz, 'GDPR reform alive again? July update' (How EU Law Influences Tech, 30 July 2026); European Digital Rights, 'The Digital Omnibus is going on summer break. Your rights are not' (16 July 2026); Joint Business Statement, 'Stop the Council Compromise on Digital Omnibus!' (June 2026).

[28] Agence Europe, 'Digital Omnibus – Irish Presidency consults EU countries on pseudonymisation, data processing for AI and cookies' (15 July 2026); Barczentewicz, 'GDPR reform alive again? July update'.

[29] Irish Presidency of the Council of the European Union, Policy Programme (June 2026), especially the commitments on competitiveness, simplification and fundamental rights.

[30] Council working-group discussion summary, 16 July 2026 (on file with the author), section I, recording Member State positions on pseudonymisation, the Article 29a vehicle and relative anonymity.

[31] Council working-group discussion summary, 16 July 2026 (on file with the author), sections II and IV, recording Member State positions on Article 88c, recital 33a, Article 9(2)(k), ePrivacy exemptions, contextual advertising, fraud prevention and audience measurement.

operative law or return them to recitals and EDPB guidance. Until that text appears, claims that the Council has rejected Article 4 reform or endorsed the Cyprus institutional architecture remain premature.

One further Council initiative sharpens this analysis. Germany and Sweden have proposed an Article 29a route that leaves Article 4(1) untouched. Still, states in the operative text that data are not personal for an entity that cannot identify the person through means that the entity possesses or can obtain and that are reasonably likely to be used. Politically, the proposal matters because two states associated with the coalition that stopped the Cyprus compromise have not rejected contextuality. Instead, they have sought a vehicle that avoids reopening the definition itself. That vehicle may offer Ireland the nucleus of a Council settlement. The location of the rule cannot, however, become a substitute for its legal effect: a Chapter IV provision must govern the GDPR's material scope across the Regulation, rather than operate only as an adjustment to controller and processor duties, and it must preserve rather than truncate the logic of SRB.[32]

As drafted, the proposal does not yet satisfy that test. Its onward-transfer clause risks making the transferor's status turn on a downstream recipient's capabilities even where identification lies outside the transferor's realistic sphere of control; its categorical processor exclusion substitutes legal status for an assessment of actual capability; and its proposed insertion of anonymisation into Article 4(2), paired with a new Article 6(1)(g) basis, could leave an express compliance burden without matched relief if the new basis falls away. The proposal also returns the operational criteria to a non-binding EDPB opinion without creating a statutory presumption or a reliable route to review. A workable compromise should confine the transfer rule to designed, known or reasonably foreseeable identification routes; attribute relevant means across a controller and processor arrangement to prevent laundering without treating the processor label as conclusive; place the objective Recital 26 factors and legal consequences in the enacting terms; make clear that anonymisation does not require consent by default; and link approved methods to a rebuttable presumption. If the co-legislators cannot secure those changes, the uncodified SRB position is preferable to legislation that recognises contextuality in principle but removes it in the cases where it matters.

Table 1.1. Commission Proposal, Council Direction and Unresolved Consequences

Issue	Commission proposal	Cyprus direction	Position on 31 July 2026
Personal data and pseudonymisation	Actor-relative Article 4 clarification plus Commission implementing power in Article 41a.	Deleted both and proposed Article 29a plus an EDPB opinion.	Article 4 remains blocked for several delegations, but Article 29a and relative anonymity now have a discernible Council coalition. The dispute concerns scope, safeguards and legal effect.
AI legitimate interests	Operative Article 88c with safeguards.	Deleted Article 88c and retained weaker recital language.	Council remains split between operative Article 88c and recital 33a. A narrower binding clarification has identifiable support.
Incidental special-category data in AI	Article 9(2)(k) derogation with Article 9(5) safeguards.	Retained in narrowed incidental and residual form.	One of the more plausible compromise provisions, subject to lifecycle controls.
Cookies and terminal equipment	Split personal-data access into the GDPR and non-personal information into ePrivacy; machine-readable choices.	Kept the trigger in Article 5(3) and retained parts of the choice-signal architecture.	A coalition supports broader low-risk exemptions for analytics, fraud prevention and contextual advertising, subject to defined safeguards and later structural review.

[32] Germany and Sweden, 'Comments on the Presidency Revised Compromise Text (P2B, GDPR, ePrivacy/Cookies, Cyber)' (joint drafting suggestions, 2 June 2026), proposed amendments to GDPR arts 4(2), 6(1)(g) and 29a; 'Digital Omnibus—Definition of Personal Data: Assessment of the SWE-GER Proposal & Advocacy Recommendation' (working analysis, July 2026), both on file with the author.

Issue	Commission proposal	Cyprus direction	Position on 31 July 2026
Automated decisions	Broader rewrite of Article 22.	Preserved the rights-based structure with limited clarification.	Parliament is likely to preserve the prohibition while debating meaningful human intervention.
Scientific research	Definition, compatibility and transparency clarifications.	Narrower definition influenced by EDPB positions.	Commercial and private research status remains politically contested.
Institutional governance	No substantial EDPB/EDPS reform; several Commission implementing roles.	Expanded reliance on EDPB guidance and opinions.	Instrument choice is now explicit: several delegations and the Commission prefer binding law or implementing acts to EDPB-only guidance.

Table 1.2. Council Political Geometry After the 16 July Working-Group Discussion

Issue	Emerging Council geometry	Strategic implication
Personal data	Article 4 remains politically blocked for several delegations, but Article 29a has become the principal compromise vehicle. Germany, Sweden and Denmark support relative anonymity; Poland and several cautious states prefer Article 29a to a definitional amendment.	Accept the vehicle only if the rule governs GDPR scope across the Regulation and preserves SRB.
Legal effect	Czechia, Poland, Hungary and the Commission favour binding criteria or implementing acts. Estonia and Cyprus remain more receptive to EDPB guidance; France seeks a fuller instrument analysis.	Place the trigger and reliance effect in law. Use EDPB methods as technical input, not as the sole source of the perimeter.
AI	Germany, Poland, Denmark, Austria, Czechia and the Commission favour an operative pathway. Belgium, the Netherlands, Malta and Cyprus prefer recital 33a; Finland and Spain remain flexible.	Draft a narrow operative Article 88c that preserves the ordinary Article 6 tests and does not create an automatic legal basis.
ePrivacy	A practical coalition supports low-risk exemptions for analytics, fraud prevention and contextual advertising. More cautious states demand narrow purposes, safeguards and impact evidence.	Separate low-risk first-party functions from cross-service tracking and build a purpose-based whitelist.
Governance	France seeks earlier and more systematic consultation before guidance; Poland has raised competitiveness in the EDPB mandate; Germany seeks further risk-based simplification.	Integrate the duty to cooperate and Innovation Mandate into the reform package rather than treating them as later administrative questions.

Parliament presents a different political problem. The joint ITRE and LIBE draft report records only the points that Aura Salla and Marina Kaljurand could place in an initial common text. It is neither a committee position nor Parliament's negotiating mandate. The committee procedure also includes IMCO and JURI opinions, a targeted impact assessment and a very large amendment exercise. Those features create several veto points, but also room for package bargaining across personal data, cookies, AI, research, enforcement, and administrative simplification.[33]

[33] European Parliament, Legislative Observatory, Procedure 2025/0360(COD), 'Simplification of the digital legislative framework – Digital Omnibus (Omnibus VII)', status accessed 31 July 2026; European Parliament, Draft Report PE786.818v01–00, 22 June 2026; European Parliament, IMCO Amendments PE791.061v01–00, 16 July 2026.

Salla's public position exposes the most useful opening. She rejects the Commission's Article 4 drafting as uncertain and vulnerable to loopholes, yet calls for a more risk-based GDPR, context-dependent treatment of anonymisation and pseudonymisation, greater use of privacy-enhancing technologies and technical standards, fewer unnecessary cookie banners and better access to data for European SMEs, research and AI development. That is not an absolute objection to contextuality. It is an argument about legal technique, safeguards and distributional effects.[34]

Kaljurand and rights-focused groups will continue to resist any text that appears to permit tracking, processor laundering or strategic corporate separation. The EDPB's July Guidelines make a compromise easier to construct because they confirm that contextuality can coexist with strict anti-circumvention rules. A processor cannot claim an independent anonymity perspective merely because it lacks the key.[35] Realistic onward transfer remains relevant, and cross-service device-fingerprint profiling remains personal data.[36] The legislative task is to place those safeguards around an operative test.

Sequencing will shape the outcome. The Presidency requested written comments by 28 July and scheduled the next AGS meeting for 11 September. Parliament is unlikely to establish its negotiating position before early 2027, while Ireland may still secure a Council mandate first. A coherent September Council text could therefore become the practical anchor for later negotiations. A second Council failure would point in the opposite direction: no durable majority would exist for meaningful reform, and the path of least resistance would favour administrative changes, recitals and supervisory guidance.

The principal political risk is an asymmetric settlement. Cookie banners, reporting templates and selected AI accommodations may offer visible compromises, while Article 4, assurance, EDPB and EDPS governance and the supervisory method remain unresolved. Such an outcome would simplify parts of the compliance surface without changing the institutions that determine the practical perimeter. It would reduce irritation while preserving the underlying source of uncertainty.

[34] Aura Salla, statement on the amendments to the Digital Omnibus report (July 2026), supporting a more risk-based GDPR, fewer unnecessary cookie banners, contextual guidance on anonymisation and pseudonymisation, privacy-enhancing technologies and access to data for research and AI development.

[35] EDPB, Guidelines 02/2026, examples distinguishing a processor acting on a controller's behalf from an independent controller such as a research institute. The processor does not obtain a separate anonymity perspective merely because it lacks the key.

[36] EDPB, Guidelines 02/2026, online-advertising example concerning pseudonyms derived from device fingerprints and used to link behaviour across services, profile users and select advertisements.

Section 2: Restoring the Risk-Based Logic of the GDPR

The GDPR established a regulatory framework capable of reconciling two foundational objectives of the European Union: protecting fundamental rights and enabling the free movement of data within the internal market.[37] These objectives are not inherently in conflict. The legislature drafted the Regulation on the assumption that effective protection of individuals' rights could coexist with a dynamic European data economy, provided that regulatory obligations reflect the risks generated by particular forms of data processing.[38] Rather than imposing uniform restrictions across all data uses, the GDPR adopted a structure in which regulatory intensity scales with the likelihood and severity of risks to individuals' rights and freedoms.[39] This risk-based orientation is not merely implicit but reflects a central organising principle of the GDPR, embedded across its provisions and legislative design.[40] As both institutional and academic analyses observe, the Regulation establishes a system in which actors calibrate compliance obligations through a structured assessment of risk to the rights and freedoms of natural persons.[41]

The Regulation's recitals explicitly reflect this risk-based orientation by requiring that the likelihood and severity of risks to individuals determine the nature and extent of compliance obligations across different processing activities. Recital 74 provides that controllers should implement appropriate measures taking into account "the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity". Recital 75 further elaborates a non-exhaustive list of risks, while Recital 76 confirms that both the likelihood and severity of potential harm must guide regulatory responses. These provisions operate through core obligations such as Article 24 (controller responsibility), Article 25 (data protection by design and by default), Article 32 (security of processing), and Article 35 (data protection impact assessments), each of which explicitly ties compliance measures to risk assessment.

The risk-based approach extends beyond isolated provisions and anchors the Regulation at a normative level. It serves as a horizontal structuring principle, linking core obligations such as accountability, security, and impact assessment to a prior risk evaluation. Both institutional and academic analyses recognise this risk-based structure as a defining feature of the GDPR. The European Data Protection Board has consistently characterised the Regulation as requiring a "risk-based approach" to compliance, particularly in the context of accountability and technical and organisational measures. Similarly, the EDPS has emphasised that the GDPR establishes a framework in which obligations must be interpreted and applied in light of the risks posed to individuals, rather than as uniform or absolute constraints on data processing.[42] Academic commentary has likewise identified risk calibration as central to the Regulation's design, noting that the GDPR "embeds risk management as a core organising principle of data protection law".[43]

This architecture is evident throughout the Regulation's design. The GDPR repeatedly emphasises that compliance obligations should reflect the risks associated with specific processing activities.[44] Controllers must assess the potential impact of their processing and implement safeguards commensurate with the identified risk.

[37] GDPR, article 1 (1), (2), (3).

[38] *Art. 29 Data Protection Working Party*, WP 218, 2014, p. 2; See also *Veil*, ZD 2015, 348 (348); *Thoma*, ZD 2013, 578 (580).

[39] See GDPR art 24(1) sentence 1, which specifies the accountability requirement by obliging the controller to implement appropriate measures taking into account the nature, scope, context and purposes of processing, as well as the risks of varying likelihood and severity to the rights and freedoms of natural persons; see also GDPR recitals 74–76.

[40] See GDPR art 5(2), read with art 24(1), requiring measures calibrated to the nature, scope, context and purposes of processing and to risks of varying likelihood and severity; GDPR recitals 74–76; Veil, 'Die Datenschutz-Grundverordnung: Überblick und erste Bewertung' (2015) ZD 348, 349.

[41] On the risk-based approach as a defining structural principle of the GDPR, see Rüpke, Lewinski and Eckhardt, Data Protection Law (2018) §18 mn 18; Lang in Taeger and Gabel (eds), GDPR/BDSG/TTDSG (4th edn, 2022) art 24 para 31; Veil, 'Die Datenschutz-Grundverordnung: Überblick und erste Bewertung' (2015) ZD 348, 349; Schuh and Weiss, ''Die Zweckbestimmung und Zweckbindung als Weichenstellung für die DSGVO-konforme Nutzung von Daten für KI-Systeme' (2024) ZfDR 225, 235. The Article 29 Data Protection Working Party had already emphasised that this approach does not constitute a novel departure in EU data protection law, but rather reflects a continuation of earlier risk-oriented elements found in Directive 95/46/EC, particularly arts 17(1) and 20. The elevation of risk calibration to a general organising principle gained particular prominence during the legislative negotiations leading to the GDPR, where it was progressively articulated as a horizontal mechanism for aligning regulatory obligations with the likelihood and severity of harm to data subjects.

[42] European Data Protection Board and European Data Protection Supervisor, Joint Opinion 01/2025 on the Proposal extending mitigating measures available for SMEs to small mid-caps and further simplification measures, 8 May 2025, para 10, describing the retained high-risk exception as consistent with the GDPR's risk-based approach (emphasis added).

[43] *Cate, Kuner, Millard, Svantesson, Lynskey*, International Data Privacy Law 2015, Vol. 5, No. 2, 95 (95 f.)

[44] Lee A Bygrave, *Data Privacy Law: An International Perspective* (OUP 2014) 168–172; Paul Voigt and Axel von dem Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide* (Springer 2017) 10–15.

Data protection impact assessments, enhanced security obligations and governance structures apply when processing is likely to create a high risk to individuals.[45] The accountability principle requires controllers not only to comply but also to demonstrate compliance in a manner proportionate to those risks.[46] This structure lets the Regulation accommodate a wide range of uses while concentrating its strongest safeguards where the potential consequences for individuals are most serious.

The advantages of such a structure are especially apparent in a regulatory environment characterised by the scale and heterogeneity of contemporary data processing. The European Commission has consistently emphasised that the digital economy depends on the ability to use data across sectors and borders, including in areas such as scientific research, industrial innovation, and the development of artificial intelligence.[47] A framework that treated all such activities as requiring identical levels of regulatory intervention would impose substantial friction across these domains. By contrast, a risk-based system allows legal obligations to adjust to different operational contexts while maintaining a consistent level of protection for individuals' rights.

Hypothetical Illustration: Uniform Treatment of Unequal Risk

Suppose a European neurological research consortium receives coded hospital data under a strict governance agreement. Each hospital retains the re-identification key; the consortium cannot access it, link the records to named patients or obtain additional identifiers through lawful means. Treating the dataset as personal data for every participant, regardless of its actual capabilities, would collapse the distinction between the originating controller and a recipient that lacks reasonable means of identification. The example is hypothetical. It isolates the legal question: whose means and circumstances determine identifiability when the recipient processes the data?

Over time, however, the practical interpretation of the GDPR has, in some areas, moved away from this calibrated model. The cited commentary describes this development as a shift toward a *de facto* "zero-risk approach" in applying the GDPR, whereby supervisory and judicial interpretations prioritise maximum protection irrespective of contextual risk differentiation. This tendency is visible, for example, in the expansive reading of concepts such as "data protection by design and by default", where guidance has, in some instances, encouraged organisations to adopt maximalist safeguards irrespective of the specific risk profile of the processing activity.[48] While precaution undoubtedly plays an important role in protecting fundamental rights, an excessively expansive understanding of risk can produce unintended consequences for the coherence of the regulatory framework. When regulatory obligations arise across a very broad range of situations, the distinction between low-risk and high-risk processing begins to erode.

This erosion has important systemic implications. Organisations navigating the regulatory environment may respond by adopting defensive compliance strategies designed primarily to minimise exposure to enforcement risk rather than to manage substantive risks to individuals. Authorities may nevertheless treat activities that pose limited practical risks as requiring extensive compliance infrastructure, documentation, and legal review. This dynamic has been observed across multiple sectors, particularly in data-driven research and AI development, where legal uncertainty over risk classification can discourage the use of otherwise valuable datasets.[49]

Over time, this process can generate a form of regulatory inflation in which the intensity of compliance obligations becomes progressively detached from the level of risk that particular processing activities actually generate. Such developments do not necessarily reflect deliberate policy choices. They often result from incremental interpretative developments across multiple institutions and enforcement contexts. Yet their

[45] GDPR arts 24, 25, 32, 35.

[46] GDPR art 5(2); See also Council of the European Union, *Council position (EU) No 6/2016 at first reading with a view to the adoption of a Regulation of the European Parliament and of the Council on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)* [2016] OJ C 159/83, 84 (emphasising that the strengthening of accountability obligations was intended to embed a risk-based approach, ensuring that controllers' duties are calibrated according to the nature, context and risks of processing and thereby fostering a 'culture of data protection').

[47] European Commission, *A European Strategy for Data* COM(2020) 66 final.

[48] EDPB, Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, version 2.0, adopted 20 October 2020, paras 6–10.

[49] European Commission, *Impact Assessment on the AI Act* SWD(2021) 84 final.

cumulative effect can alter how a regulatory framework functions in practice. A system originally designed to calibrate obligations according to risk may gradually come to operate as though most processing activities require the same level of precautionary scrutiny. When this occurs, the internal architecture of the regulatory framework begins to lose the proportional structure that originally made it workable.

The stakeholder submissions reviewed for this report describe the same mechanism in practical terms. La Villa Numeris reported that the GDPR had strengthened trust but had not dispelled a climate of anxiety, divergent interpretation or disappointment with cumbersome procedures and the limited practical uptake of tools such as codes of conduct and certification. Its proposed response was not abandonment of protection, but a change in regulatory culture: more constructive soft law, proportionate enforcement and stronger support for organisations.[50]

Commentary on *EDPS v SRB* further develops the point. It argues that treating pseudonymised data as personal whenever re-identification is theoretically possible encourages fear rather than responsibility. At the same time, the June 2026 intervention stresses that differing national readings remain a barrier for pan-European actors. These are advocacy sources rather than causal studies. Their value lies in showing how regulated actors experience uncertainty: as an incentive to adopt the strictest plausible position, to delay sharing, and to relocate product decisions away from the legal margin.[51]

What This Illustrates

When risk is interpreted expansively without clear differentiation:

- Low-risk and high-risk processing are treated alike
- Compliance decisions prioritise enforcement avoidance over safeguards
- Research and innovation activity is constrained or displaced
- Regulatory attention is misallocated

A risk-based framework does not reduce protection. It ensures that protection is **targeted, effective, and sustainable.**

Restoring the intended functioning of the GDPR, therefore, requires renewed emphasis on three interrelated regulatory principles: ***legal certainty***, ***risk-based governance***, and ***proportionality***. These principles do not sit outside the framework. Still, they are embedded within the structure of EU law itself, including the Charter of Fundamental Rights and the general principles governing Union legislation.[52] The challenge lies not in identifying these principles but in ensuring that they continue to guide the interpretation and development of the regulatory framework in practice.

Legal certainty is a foundational element of the rule of law within the European Union. It requires that legal rules be clear, predictable, and consistently applied.[53] Within the context of the GDPR, this principle is particularly important because the Regulation operates on a trigger-based system. Once a dataset or activity falls within the scope of personal data processing, a complex set of obligations becomes applicable. Uncertainty about the boundaries of that trigger condition therefore propagates throughout the regulatory system, affecting compliance decisions across a wide range of economic and institutional contexts. EU jurisprudence explicitly recognises this dynamic, as the Court of Justice emphasises the need for clear and precise rules when fundamental rights are engaged.[54]

Risk-based governance provides the mechanism through which the GDPR allocates regulatory attention to areas where it is most needed. By linking compliance obligations to the level of risk generated by particular processing activities, the Regulation enables regulators and organisations alike to focus resources on areas where

[50] La Villa Numeris, GDPR: 5 years, is that something to celebrate? (Report summary, 25 May 2023), identifying divergent interpretation, compliance anxiety and underused accountability tools, and proposing a more constructive approach to soft law and regulatory support.
[51] David Lacombled, 'GDPR: The Return of Common Sense' (L'Opinion, 23 September 2025); David Lacombled, 'Europe: Simplifying the General Data Protection Regulation without betraying it' (L'Opinion, 2 June 2026).
[52] Charter of Fundamental Rights of the European Union, arts 7, 8, 52.
[53] Case C-17/03 *VEMW* EU:C:2005:362.
[54] Joined Cases C-293/12 and C-594/12 *Digital Rights Ireland* EU:C:2014:238.

individuals' rights are most likely to be affected. This approach is not unique to the GDPR but reflects a broader shift in EU regulatory practice toward risk-based governance models, particularly in complex technological domains. Within the GDPR itself, this principle is reinforced through the requirement that both controllers and supervisory authorities assess risk dynamically, rather than relying on static or categorical classifications of data processing activities.

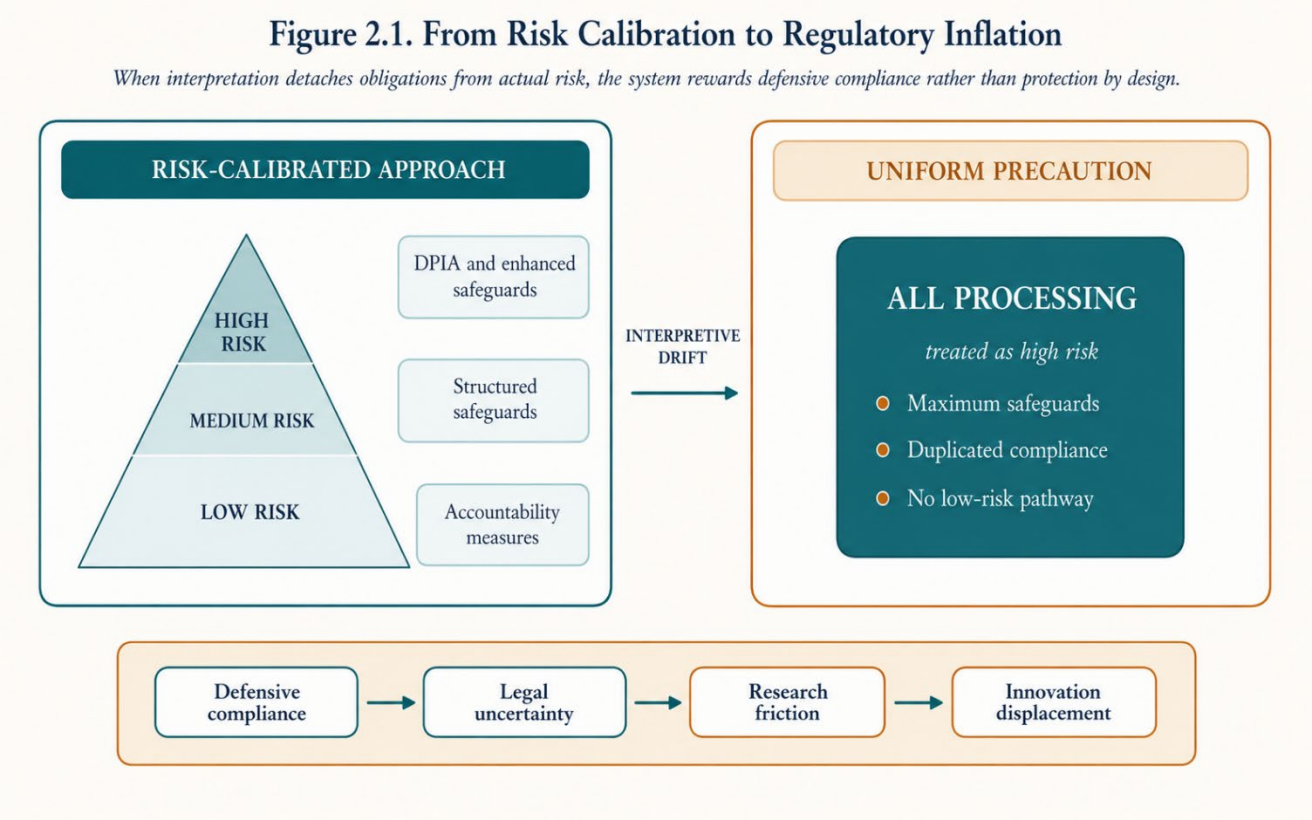

The principle of ***proportionality*** complements this risk-based approach by ensuring that regulatory measures remain appropriate and necessary to achieve their objectives. Within EU law, proportionality operates as a general constraint on regulatory action, requiring that measures do not impose excessive burdens relative to the aims pursued.[55] Within the GDPR framework, this principle ensures that data protection obligations remain compatible with other objectives of Union law, including scientific research, technological development, and the functioning of the internal market.[56] Without such balancing, the regulatory framework risks generating friction with other components of the EU legal order, particularly as new digital legislation expands the scope of regulatory intervention.

More recent components of the EU digital regulatory framework also reflect the importance of calibrated regulatory governance. The Artificial Intelligence Act organises obligations around differentiated levels of regulatory intensity, with obligations increasing in line with the risks associated with specific categories of AI systems.[57] The Commission has framed this approach as necessary to protect fundamental rights and promote innovation within the European AI ecosystem.[58] Although the AI Act and the GDPR regulate different technological and legal phenomena, both instruments embody a common legislative preference for risk-tiered governance rather than uniform restriction.

This broader pattern across the EU digital *acquis* reinforces the importance of preserving the internal logic of the GDPR's risk-based structure. As the regulatory landscape expands to include new instruments governing AI, data sharing, and digital platforms, coherence between these frameworks becomes increasingly important. If the

[55] Consolidated Version of the Treaty on European Union, art 5(4); Charter of Fundamental Rights of the European Union, art 52(1); *Case C-331/88 R v Minister of Agriculture*, Fisheries and Food, ex p Fedesa EU:C:1990:391, para 13.
[56] GDPR recitals 4 and 10 and arts 1(1)–1(3), linking protection of fundamental rights with the free movement of personal data within the Union.
[57] Regulation (EU) 2024/1689 (Artificial Intelligence Act), arts 6, 9, 10, 14 and 26, establishing risk-tiered duties for high-risk AI systems and deployers.
[58] European Commission, Proposal for a Regulation laying down harmonised rules on artificial intelligence, COM(2021) 206 final, explanatory memorandum, 1–3; Artificial Intelligence Act recitals 1 and 26.

GDPR were to drift toward a model in which regulators treat most data processing activities as inherently high risk, tensions with other parts of the digital regulatory framework would likely intensify, particularly in areas such as AI training, data sharing, and cross-sector innovation.

The Omnibus can restore the Regulation's risk-based architecture only if it addresses both legal text and supervisory method. Clear provisions matter, but open-textured standards will continue to require guidance, technical assessment and enforcement choices. Without procedural disciplines at that administrative layer, narrow drafting gains can be lost through precautionary interpretation, inconsistent evidence thresholds and remedies that treat the most restrictive lawful route as the default.

A bounded Innovation Mandate supplies the missing complement. It does not create a defence to unlawful processing or place growth above Articles 7 and 8 of the Charter. It applies only where legislation leaves more than one lawful route. In that choice set, a regulator should identify the rights or safety risk, disclose the evidence and assumptions, test less burdensome rights-preserving alternatives, consider effects on legal certainty, SMEs, investment and scale, and explain why the selected approach remains proportionate. The mandate disciplines method rather than predetermining outcomes.[59]

2.1 Transition: Why Institutional Balance and Supervisory Method Matter

The simplification debate therefore raises a question of institutional allocation as well as substantive law. The EDPB, the EDPS and national supervisory authorities perform indispensable interpretative, coordinating and enforcement functions. Their decisions and guidance increasingly shape conduct across the GDPR, AI Act, DSA, DMA, Data Act, ePrivacy and sectoral data spaces. That expanded influence requires clearer legal boundaries, stronger cooperation duties and more visible administrative reasoning.

Guidance can reduce uncertainty, but it can also generate quasi-legislative expectations when authorities, auditors, engineers and investors treat it as the practical benchmark for later enforcement. The distinction between binding and non-binding instruments remains legally important; it does not describe their full market effect. Where the basic act leaves the perimeter open, the least contestable legal form may become the most influential one.

The constitutional concern does not arise because supervisors protect rights vigorously. It arises when public bodies make scope-affecting choices without procedures proportionate to their practical authority. Independence protects decisional integrity. It does not remove the need for evidence, consultation, reasons, recorded disagreement, cross-regulatory cooperation and judicially intelligible responsibility.

Where legislation fails to settle the trigger, four consequences follow.

First, guidance and consistency opinions become the practical source of perimeter rules.
Second, organisations anticipate the strictest plausible interpretation because misclassification carries high penalties and sunk design costs.
Third, national divergences survive beneath a nominally common EDPB position.
Fourth, courts receive the dispute only after investment, product design or research collaboration has already been altered.

The resulting system rewards defensive over-compliance rather than demonstrable risk reduction. It also distributes cost unevenly: incumbents can purchase specialist interpretation and run parallel architectures, while smaller firms and research bodies often cannot.

The Omnibus should therefore restore a clear division of labour. The co-legislators must define the perimeter and the legal consequences of satisfying an approved assurance route. The EDPB should supply evidence, methodologies, technical updates, codes and consistency. The EDPS and national authorities should enforce within that framework. Cross-regulatory bodies should cooperate under explicit, reviewable duties. Courts should remain able to identify which institution made each determinative choice.

[59] Leiser, 'The Innovation Mandate', proposing a bounded secondary duty that applies only within a lawful choice set and requires evidence, alternatives, proportionality and public reasons for major supervisory positions.

Section 3: Institutional Balance — Legislative Settlement, Supervisory Authority and the Innovation Mandate

The central institutional question is who determines the GDPR's practical perimeter. The answer carries constitutional and economic consequences because Article 4 activates a dense system of duties, rights and penalties and now controls how several adjacent digital regimes interact. Concepts such as identifiability, reasonable means, role allocation and effective anonymisation require judgment. The legislature must decide which parts of that judgment constitute essential elements and which parts can properly remain technical or case-specific.

The GDPR gives the EDPB a central role in consistency, guidance and dispute resolution; it does not make the Board a co-legislator. Guidelines, recommendations and Article 64 opinions coordinate interpretation but remain formally non-binding. Article 65 decisions bind supervisory authorities in defined cross-border disputes and may directly affect regulated parties. The legal forms differ because they carry different procedures, effects and routes of review.

Soft law complicates the distinction. Organisations often follow EDPB guidance because later departure carries enforcement, audit and redesign risk. National authorities cite it, courts consider it and technical standards absorb it. The market therefore experiences much of the discipline of binding law even where direct annulment review remains unavailable. That gap between formal force and practical effect makes *ex ante* procedure especially important.[60]

Perimeter drift emerges when guidance does more than explain a legislated threshold and instead supplies the threshold's decisive content. Each move may appear modest. Together, supervisory expectations, anticipatory compliance and enforcement feedback can widen the practical scope without any amendment to Article 4.

Figure 3.1.1. Directional Drift in the Practical Scope of the GDPR

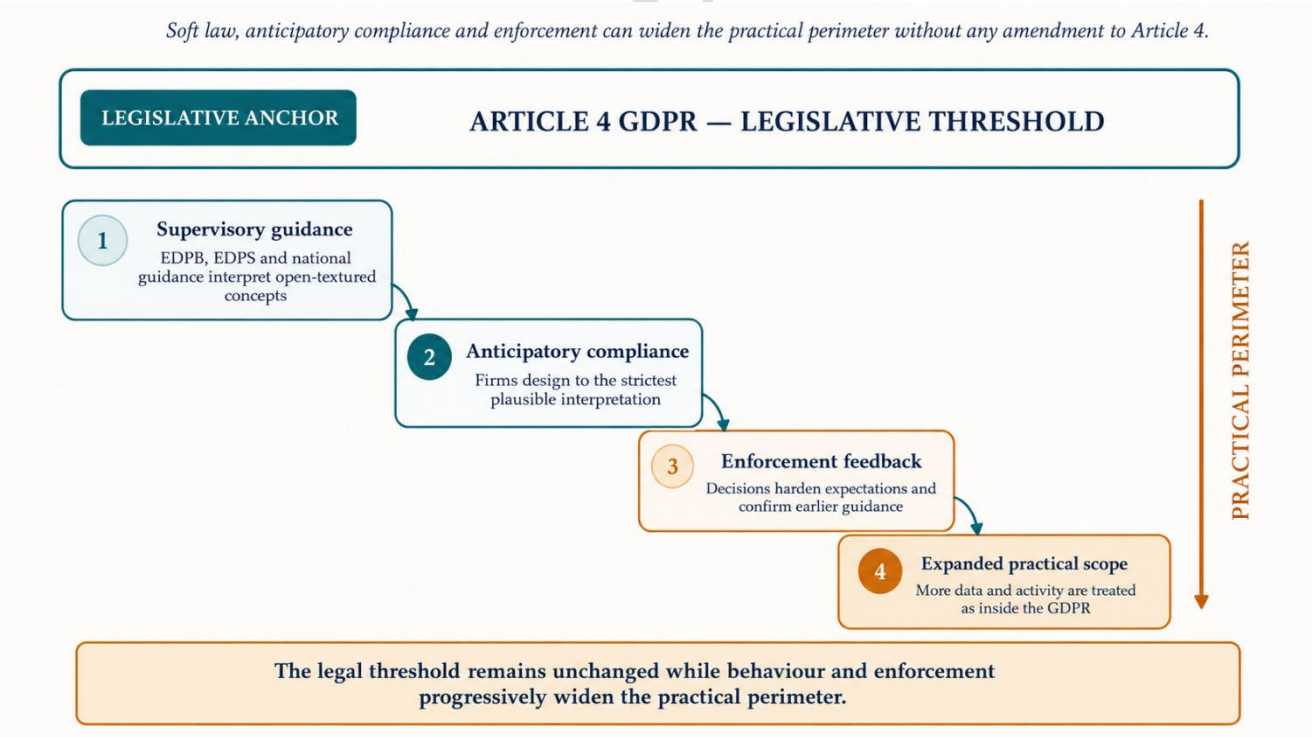

Figure 3.1.1 illustrates this directional effect. Organisations do not respond to abstract legal doctrine alone. They respond to the strictest credible enforcement position, the cost of redesign and the possibility that a later authority will reject a contextual classification. The result is a one-way ratchet unless legislation supplies a reliable path for demonstrating that risk-reducing architecture changes legal status.

[60] *Case C-322/88 Grimaldi v Fonds des maladies professionnelles EU*:C:1989:646; *Case C-16/16 P Belgium v Commission EU*:C:2018:79; Case C-911/19 *Fédération bancaire française v Autorité de contrôle prudentiel et de résolution* EU:C:2021:599.

> Two features intensify the ratchet: the GDPR relies on open-textured, risk-based language, and its penalties and remedial exposure make *ex post* correction commercially inadequate. A nominally flexible standard can become more restrictive than a clear prohibition when actors cannot know in advance which evidence will satisfy it.

This dynamic raises more than a compliance-cost question. It concerns the democratic allocation of market-shaping power. Digital supervisors translate broad statutes into guidance, priorities, templates, remedies, sandbox criteria and technical expectations. Those choices allocate delay, investment and opportunity. The Union should therefore govern them through procedures that expose evidence, alternatives, distributional effects and proportionality rather than treating implementation as a neutral administrative afterthought.[61]

The Court's judgment in *WhatsApp Ireland v EDPB* supplies one accountability anchor. It confirms that an Article 65 decision may constitute a challengeable Union act where it determines the legal position that the lead authority must implement. The judgment does not make all EDPB guidance reviewable. It does, however, reject the idea that the Board can exercise binding public authority while presenting the result as merely internal coordination.[62]

EDPS v SRB supplies the substantive anchor. The Court rejected the proposition that pseudonymised information must remain personal data for every person in every context. It required an actor-specific assessment of means reasonably likely to be used, while preserving responsibility where disclosure creates a realistic identification route. That combination makes contextuality compatible with protection; it also makes legislative precision indispensable.[63]

The EDPB's response has now moved from preparation to an adopted consultation text. Guidelines 02/2026 state that the anonymity of information may differ between entities, require assessment from each relevant entity's perspective and describe the contextual approach as the full legal standard. They also retain a simplified approach that treats information as personal whenever it was personal for another entity or arose from personal data, even though that approach may over-include information that is anonymous in the relevant context.[64]

That creates an institutional dilemma. In February, the EDPB and EDPS argued that the co-legislators should delete the Article 4 amendment and rely on forthcoming guidance.[65] In July, the promised guidance accepted the core contextual principle that civil-society opponents had characterised as a deregulatory novelty.[66] If guidance was authoritative enough to justify legislative inaction, its contextual conclusion must carry weight. If guidance remains too uncertain or non-binding to settle the issue, the case for legislation strengthens.

A proportionate response is neither to ignore the Guidelines nor to convert them into the sole source of the perimeter. Their anti-circumvention rules, operational criteria and examples should inform the basic act. Their remaining ambiguity, absence of reliance consequences and capacity for revision show why guidance cannot replace it.

The Digital Omnibus now presents four institutional models rather than two. The Commission proposed legislative clarification plus a Commission implementing power. Cyprus removed both and relied more heavily on EDPB opinion and guidance. Germany and Sweden have proposed an operative Article 29a rule that accepts entity-relative classification but surrounds it with broad transfer and processor exclusions, while still leaving technical criteria to the Board. The model argued for here places the trigger, burden, anti-circumvention rules and legal consequence in the basic act, while assigning technical methods, testing and updates to the EDPB within that fixed perimeter.

---

[61] M R Leiser, 'The Innovation Mandate: Making Europe's Digital Supervisory State Accountable for Rights, Growth and Scale' (working paper, July 2026) <https://osf.io/preprints/socarxiv/bmcnv_v1> accessed 31 July 2026.

[62] *Case C-97/23 P WhatsApp Ireland Ltd v European Data Protection Board EU*:C:2026:81.

[63] *Case C-413/23 P European Data Protection Supervisor v Single Resolution Board EU*:C:2025:645, especially paras 82–100.

[64] EDPB, Guidelines 02/2026, which distinguish a contextual approach reflecting the full legal standard from a simplified approach that may treat information as personal even where it is anonymous in the relevant context.

[65] European Data Protection Board and European Data Protection Supervisor, Joint Opinion 2/2026 on the Proposal for a Regulation as regards the simplification of the digital legislative framework, adopted 10 February 2026, especially paras 13–25 and 96–117.

[66] European Data Protection Board, Guidelines 02/2026 on Anonymisation, version 1.0, adopted for public consultation on 7 July 2026, especially the Executive Summary and sections 2.2, 3 and 4.

The Commission's Article 41a raised a legitimate concern about essential elements. Criteria that decide whether information falls outside the GDPR can alter the regime's material scope. An implementing act may standardise methods and technical evidence, but it should not possess an open-ended power to decide the legal trigger. The correct response is to place the trigger, burden, anti-circumvention rules and reliance effect in primary legislation and reserve implementation for auditable detail.[67]

The EDPB and EDPS Joint Opinion correctly identified institutional risk, but its proposed remedy left the opposite risk untouched. Assigning the question to supervisory interpretation does not remove perimeter-setting power from the administrative layer. It relocates that power to instruments with weaker democratic settlement and, in the case of guidelines and Article 64 opinions, weaker direct reviewability.

The July Guidelines make the allocation problem concrete. They contain valuable legal and technical analysis, but they do not create a statutory presumption, a certification-linked consequence or a uniform defence for reasonable reliance. An organisation may follow their methodology and still face a later authority that applies the simplified approach, disputes the auxiliary-data assessment or identifies a new route to re-identification. That is a reason to require continuing review, not a reason to deny all *ex ante* legal effect.

The 16 July discussion also undermines the proposition that EDPB guidance offers the natural compromise. The Commission described implementing acts as the more robust harmonisation instrument because they remain subject to comitology and judicial review; Poland, Czechia and Hungary likewise demanded binding criteria or doubted guidance alone. Sweden warned that overly complex guidance could defeat simplification. The emerging Council question is therefore not whether technical expertise is needed. It is whether technical expertise should supply evidence within a legislative settlement or determine the perimeter through instruments that regulated actors cannot reliably invoke or challenge.[68]

3.1 from Article 41a to Article 29a, the German-Swedish Route and the July Guidelines

The Cyprus compromise replaced Article 41a with a proposed Article 29a that would restate the 'means reasonably likely' test and request an EDPB Article 64(2) opinion on pseudonymisation and anonymisation. That text never became a Council mandate. The subsequent German-Swedish proposal makes Article 29a more substantive by placing an entity-relative rule in the enacting terms, but its transfer carve-out, processor exclusion and reliance on a Board opinion risk narrowing the very principle it seeks to codify. More importantly, the EDPB has already issued Guidelines 02/2026 before the Omnibus was agreed, showing that the Board did not need Article 29a to develop a methodology.

The early guidance is useful but exposes Article 29a's weakness. An Article 64(2) opinion or guideline can coordinate authorities and influence practice; it does not bind courts or create a general rule of scope. Nor does it tell a controller what legal consequence follows from meeting an approved technical test. The problem is therefore not a shortage of guidance. It is the absence of a reliance-worthy legislative bridge between evidence and classification.

Legal certainty requires an *ex-ante* rule that organisations can use before transferring data, financing a research project or deploying a privacy-preserving product. That rule should state whose capabilities count, how access through processors, affiliates and third parties is treated, when a contemplated onward transfer changes the analysis and what evidence shifts the burden. Guidance can populate those categories with examples; it cannot supply their democratic authority.

The Board's network composition makes that distinction especially important. It brings together national supervisory authorities and the EDPS, but national legal cultures, enforcement priorities and judicial routes remain. A common text can reduce divergence without eliminating it. Legislative harmonisation and regulatory coordination perform different functions.

[67] Articles 290 and 291 TFEU and the settled distinction between essential legislative elements and measures that ensure uniform conditions for implementation.

[68] Council working-group discussion summary, 16 July 2026 (on file with the author), including the Commission's preference for implementing acts subject to comitology and judicial scrutiny and the interventions of Czechia, Poland, France, Hungary and Sweden on the legal effect of EDPB guidance.

Neither Article 29a nor Guidelines 02/2026 provides a rebuttable presumption for a documented architecture. The controller remains responsible, but the law offers no corresponding protection for reasonable reliance on an approved methodology. High penalty exposure therefore continues to favour the strictest classification even where the actor has removed realistic identification capability.

Technical change also requires a distinction between a stable legal test and an updating evidential layer. Attack methods, auxiliary datasets and computing resources evolve. The EDPB should update methodologies and safety margins accordingly. The legal consequence of satisfying the statutory standard should not change silently whenever a guidance document changes.

The contrast with Article 65 remains instructive. A binding EDPB decision can be challenged because it determines the outcome of a composite enforcement procedure. Guidance may shape the same investment and design choices while remaining harder to contest directly. Foundational scope questions should not be placed in the legal form that combines strong practical influence with the weakest reliance and review architecture.

Table 3.1. Why Guidelines 02/2026 Do Not Cure the Legal-Certainty Problem

Problem	Guidelines 02/2026 / Article 29a response	Remaining gap
Legal force	Guidelines and an Article 64(2) opinion coordinate interpretation.	Neither creates a generally binding perimeter rule.
Contextuality	The Guidelines accept entity-relative anonymity and the contextual legal standard.	Article 4 still lacks an operative formulation and anti-circumvention rules.
Reliance	Controllers remain responsible and may follow the methodology.	No statutory presumption, assurance consequence or defence for reasonable reliance.
National divergence	Common guidance can narrow disagreement.	DPAs and courts may still apply open-textured concepts differently.
Technical change	Guidance can update attack models and methods.	The legal threshold should not move silently with each guidance revision.
Reviewability	Soft law may influence practice and later enforcement.	Direct challenge remains weaker than for an Article 65 binding decision.

3.2 Reconstructing the EDPB: Technical Authority Inside a Legislatively Fixed Perimeter

Rejecting perimeter substitution does not marginalise the EDPB. It assigns the Board the functions for which a transnational expert network has comparative advantage. The legislature should determine the trigger, burden, anti-circumvention rules, circulation consequences and legal effect of approved assurance. The Board should translate that settlement into methods that remain technically current and consistent across sectors.

The EDPB can maintain a common taxonomy of attack models, auxiliary datasets, legal access routes, temporal horizons and environmental controls. It can specify how No Record Isolation, No Linkage and No Inference operate across recipient classes; organise independent testing; maintain sector examples; and state which changes require reassessment. The July Guidelines supply the beginnings of that technical framework.[69]

Codes of conduct and certification should carry a defined legal consequence. Where an actor follows an approved methodology, documents the environment and obtains independent assurance, the GDPR should create a rebuttable presumption that the classification was reasonable at the assessment date. The presumption should fail where facts were concealed, organisational separation was artificial, re-identification was procured, onward transfer escaped the assessed boundary or capabilities materially changed.[70]

Consistency also requires a public evidential record. The Board should maintain a searchable register of significant anonymisation and pseudonymisation decisions, publish divergence reports and explain why new

[69] EDPB, Guidelines 02/2026, sections on No Record Isolation, No Linkage and No Inference, together with the discussion of realistic means, auxiliary information, contracts, onward transfer and periodic reassessment.

[70] GDPR arts 40–43 on codes of conduct and certification, read with the accountability principle in arts 5(2) and 24.

technical evidence changes an earlier method. National authorities should remain independent, but departures from common methodology should be reasoned and visible.

Binding authority should remain tied to procedures that carry rights of defence and judicial review. Article 65 decisions may resolve defined disputes. General guidance should support actors and authorities rather than acquire unreviewable quasi-legislative status. The governing division of labour remains straightforward: legislation sets the perimeter; the EDPB supplies evidence, method, assurance and consistency; authorities enforce; courts review.

Table 3.2. A Reconstructed Role for the EDPB

EDPB or EDPS function	Appropriate instrument	Institutional limit
Evidence hub	Common taxonomy of attacks, auxiliary data and realistic capabilities.	Does not alter the statutory meaning of identifiability.
Methodology setter	Assessment templates, test protocols, sector examples and reassessment triggers.	Methods remain anchored in Article 4 and Recital 26.
Assurance coordinator	Codes and certification under Articles 40–43.	Creates rebuttable assurance, not immunity.
Consistency monitor	Decision register, divergence reports and reasoned comparisons.	National authorities remain independent; departures become visible.
Binding adjudicator	Article 65 dispute resolution in defined cross-border cases.	Binding acts remain case-specific and reviewable.
Cooperation duty	Notice, lawful exchange, reasons for divergence and remedy coordination.	No consensus requirement and no inter-regulatory veto.
Innovation Mandate	Impact statement, evidence register, alternatives and proportionality reasons for major general positions.	Operates only inside the lawful choice set and never excuses breach.

Cross-regulatory governance must also become a legal discipline rather than an informal aspiration. The same data architecture may engage the GDPR, AI Act, DSA, DMA, Data Act, consumer law, competition law and cybersecurity rules. A privacy authority's classification can therefore alter another regulator's mandate or remedy. The EDPB's July template for cooperation agreements recognises the coordination problem, but a template does not create a horizontal duty or determine the consequences of serious non-cooperation.[71]

The Omnibus should treat institutional balance as a design objective. A workable perimeter must be legislatively legible; supervisory governance must remain independent but procedurally disciplined; and affected actors must be able to identify which public body supplied the evidence, interpretation and remedy that determined the outcome.

1. Primary legislation should codify actor-relative identifiability, realistic indirect access, disclosure responsibility and the burden of demonstrating non-identifiability. The basic act should also define the legal effect of compliance with an approved assurance route.
2. Implementing acts should standardise formats, technical protocols and interoperable evidence where uniform conditions are necessary. They should not carry an open-ended power to redefine what falls inside or outside the GDPR.
3. Major EDPB and EDPS guidance should include a clear problem definition, an evidence note, a consultation report, a summary of material disagreement, a cross-regulatory compatibility assessment and a review date. The institutions should distinguish legal interpretation from policy choice and explain why rejected, less burdensome rights-preserving alternatives were inadequate.[72]

[71] European Data Protection Board, Template for Cross-Regulatory Cooperation Agreements, adopted 7 July 2026.

[72] M R Leiser, 'The Innovation Mandate: Making Europe's Digital Supervisory State Accountable for Rights, Growth and Scale' (working paper, July 2026) <https://osf.io/preprints/socarxiv/bmcnv_v1> accessed 31 July 2026.

4. A legally enforceable duty to cooperate should apply when one digital regulator materially assesses conduct governed by another regime, when parallel proceedings risk incompatible findings or cumulative sanctions, or when guidance will shape concepts used by another authority. Cooperation should require notice, lawful information exchange, time-bounded responses, protection of confidentiality and defence rights, reasons for material divergence, and coordination of remedies, without creating an inter-regulatory veto.[73]
5. The EDPB also needs a more autonomous administrative base or audited statutory separation from the EDPS. The EDPS supervises Union institutions, advises the legislature, supplies the EDPB Secretariat and performs AI Act functions. Independence in decision-making does not answer questions about staff, access, legal advice, records, budget and institutional memory. Clear chains of responsibility strengthen rather than weaken both bodies.[74]

An Innovation Mandate should apply to major, prospective and generally significant supervisory choices, not to individual complaints or enforcement outcomes. Within the lawful choice set, it should require regulators to disclose the risk, evidence, alternatives, cross-border effects, SME and scale-up consequences and proportionality reasons. The rights floor remains absolute; the method of choosing above that floor becomes accountable.[75]

[73] Leiser, 'Recalibrating the EDPB and the EDPS', proposing an enforceable and reviewable duty to cooperate, procedural discipline for major guidance, clearer administrative separation and auditable information governance.

[74] M R Leiser, 'Recalibrating the EDPB and the EDPS Under a Duty to Cooperate: An Institutional Case for Cooperative Digital Administration' (policy paper, July 2026) <https://osf.io/preprints/lawarchive/4mxz6_v1> accessed 31 July 2026.

[75] Leiser, 'The Innovation Mandate', proposing a bounded secondary duty that applies only within a lawful choice set and requires evidence, alternatives, proportionality and public reasons for major supervisory positions.

Table 3.3. Comparative Table of Perimeter-Relevant Instruments

Instrument Type	Legal Force and Reviewability	Typical Process and Institutional Author	Typical Effect on Scope
Primary Legislation (EU Regulations and Directives)	Binding. Reviewable under Article 263 TFEU by privileged applicants; validity questions may arise through Article 267 TFEU.	Ordinary legislative procedure, political negotiation and amendment. Institutional author: EU co-legislators acting on a Commission proposal.	Defines the perimeter and essential elements; sets triggers, rights and duties.
Implementing Acts (Article 291 TFEU)	Binding. Reviewable under Article 263 TFEU and capable of indirect challenge in national litigation.	Commission adoption through comitology, with examination or advisory procedures and technical specification. Institutional author: Commission, and exceptionally the Council.	Standardises implementation and may affect the perimeter where used to specify definitional thresholds.
EDPB Guidance (Article 70 GDPR)	Non-binding soft law. Usually not directly challengeable as such, but it may shape interpretation and be relied upon in disputes.	Drafted through the Secretariat and expert subgroups, often followed by public consultation and plenary adoption. Institutional author: EDPB.	Shapes practical scope through interpretative expectations, compliance templates and "reasonable means" methodologies.
EDPB Binding Decisions (Article 65 GDPR)	Binding on supervisory authorities and determinative for the regulated party in the composite procedure. A challengeable act where the standing requirements are met.	Triggered by reasoned objections in cross-border cases; the Board resolves the dispute, and the lead authority issues the final decision. Institutional author: EDPB.	Produces strong, case-specific perimeter effects that may generalise through precedent and enforcement signalling.
National DPA Guidance and Enforcement Positions	Guidance is generally non-binding; enforcement decisions bind under national administrative law. Review routes vary, but enforcement decisions are usually appealable nationally.	National consultation and enforcement practices vary. Institutional author: national supervisory authority.	May expand or narrow the practical perimeter domestically; divergence drives cross-border uncertainty and highest-standard compliance.

The decisive distinction is not between 'regulation' and 'innovation'. It is between enacted law and an administrative layer that can shape market conditions without equivalent procedural accountability. The Omnibus should correct that mismatch while the legislative window remains open.

3.3 The Procedural Settlement: A Duty to Cooperate and an Innovation Mandate

Substantive reform will underperform unless the Omnibus also governs the institutions that translate open legal standards into market conditions. The EDPB and EDPS now influence privacy, AI, platform governance, competition, consumer protection, cybersecurity and data access. Their practical authority has grown faster than the procedural discipline surrounding major guidance, information exchange and cross-regulatory disagreement.

A duty to cooperate should mean reasoned non-unilateralism. An authority should identify material overlap early, exchange necessary information on a lawful basis, respond within a defined period, protect confidentiality and defence rights, explain material divergence and coordinate remedies where cumulative effects matter.

Cooperation should not require consensus or create a veto; it should make disagreement visible and reviewable.[76]

The duty should bind the EDPB and EDPS themselves, not merely decorate relationships between national authorities. It should also require auditable separation, with the EDPS supplying the EDPB Secretariat while exercising supervisory, legislative-advisory, and AI Act functions. Clear records, access controls, reporting lines and responsibility for legal advice are ordinary requirements of legitimate administration, not threats to independence.

The Innovation Mandate addresses a related gap. When a supervisor chooses between lawful routes with materially different burden and market effects, the authority should publish the risk, evidence, assumptions, alternatives, and proportionality reasons. The mandate should exclude individual complaints, investigations and sanctions. It should attach to general guidance, templates, priorities, sandbox rules, consistency positions, and cross-regulatory arrangements.[77]

Together, the two reforms create a constitutionally disciplined supervisory state: cooperation allocates responsibility across institutions, while the Innovation Mandate disciplines discretion within each institution's lawful choice set. Neither reform weakens the rights floor. Both make the exercise of public power more intelligible and contestable.

Figure 3.1.2. From the EDPB's Contextual Concession to a Legislative Settlement

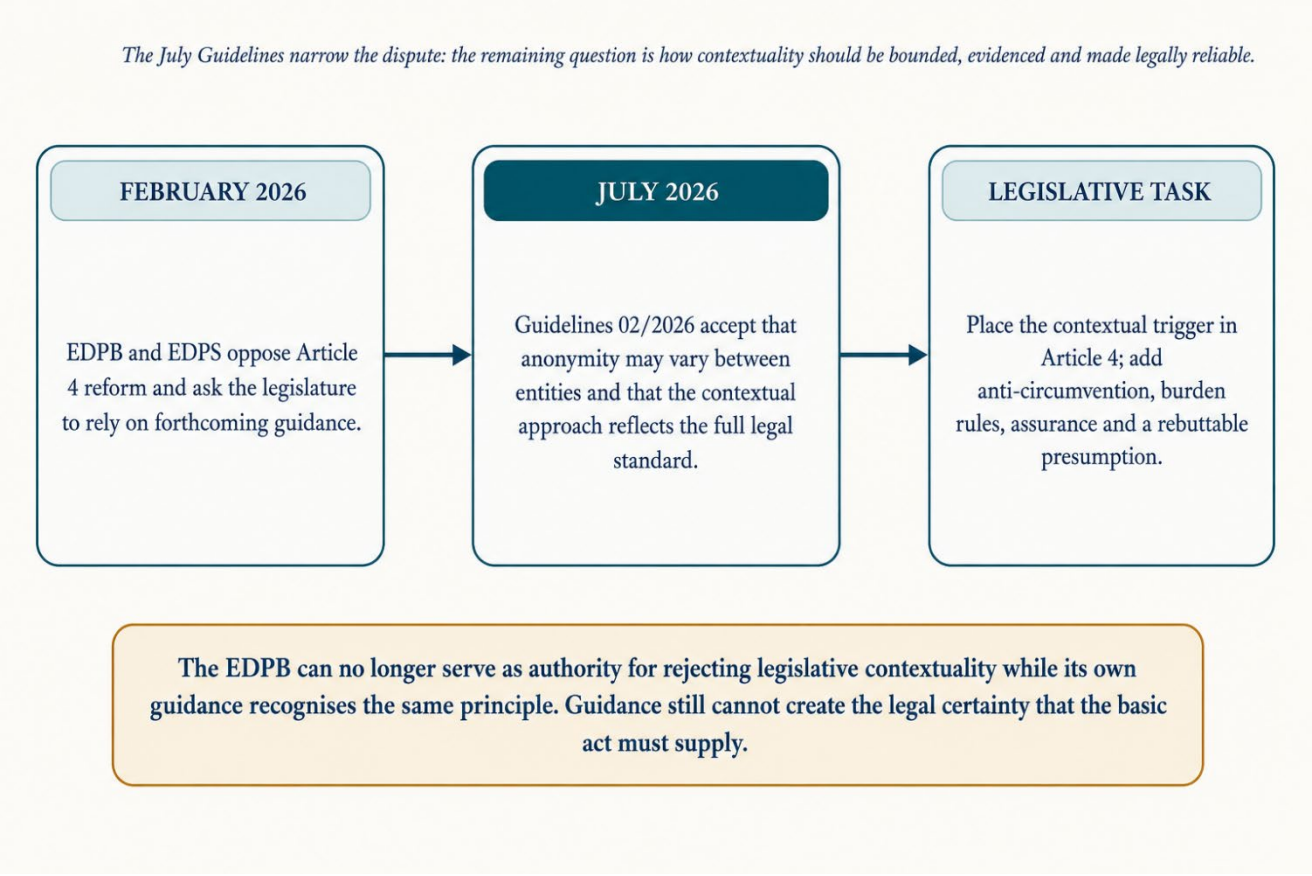

The EDPB's July concession narrows the dispute from whether contextuality exists to how legislation should bound and operationalise it.

[76] Leiser, 'Recalibrating the EDPB and the EDPS', proposing an enforceable and reviewable duty to cooperate, procedural discipline for major guidance, clearer administrative separation and auditable information governance.
[77] Leiser, 'The Innovation Mandate', proposing a bounded secondary duty that applies only within a lawful choice set and requires evidence, alternatives, proportionality and public reasons for major supervisory positions.

Section 4: The SRB Judgment and the Constitutional Question of Trigger Logic

The Court of Justice's judgment in *EDPS v SRB* clarified an important aspect of identifiability under the GDPR. The Court recognised that whether data are 'personal' depends not only on the data themselves but on the context in which an actor holds them and the actor assessing them. Pseudonymised information may constitute personal data for one actor who possesses the means to re-identify individuals. In contrast, for another actor lacking such means, the same dataset may fall outside the Regulation's practical scope. This reasoning directly affects the GDPR's trigger logic, as the definition of personal data determines when the entire regulatory framework applies.

The judgment places Article 4 at the centre of the Regulation's architecture. The definition of personal data does not merely describe scope; it determines when obligations arise and how far they extend across different actors and contexts. Identifiability, therefore, serves as a threshold condition that shapes the GDPR's reach in practice. Interpreted coherently, it allows the Regulation to protect fundamental rights without extending to all forms of data use. Article 4 also anchors the GDPR's relationship with the wider EU digital and data *acquis*. Where interpretation expands the definition without clear limits, it extends the Regulation across the digital environment beyond its intended scope. The Omnibus reform, therefore, requires a legislative decision on how that threshold should operate in practice.

A coherent interpretation preserves that calibration, allowing the Regulation to protect fundamental rights without collapsing into an undifferentiated system of control over all data processing. It is also a cornerstone provision for the relationship between the GDPR and the wider EU digital and data *acquis*. Where interpretation expands the definition without clear limits, it turns the GDPR into a law governing all forms of data use across the digital environment, rather than a framework that applies only under specific conditions. **The Omnibus reform, therefore, presents the legislature with a constitutional choice about how the trigger logic of EU data protection law should operate:**

Option 1 – Clarify Contextual Identifiability While Preserving the Definition of Personal Data

- Codify the contextual logic of identifiability recognised by the Court in *SRB*.
- Preserve the existing definition of personal data as the structural foundation of the GDPR.
- Reinforce the Regulation's risk-based and proportionate regulatory architecture. • Stabilise the role of the GDPR within the broader EU digital and data *acquis*.

Option 2 – Leave the Definition Unchanged Without Clarification

- The practical meaning of identifiability develops through guidance and enforcement practice.
- Supervisory interpretation increasingly determines the scope of the GDPR.
- Divergent national interpretations emerge across Member States.
- The GDPR progressively expands into a general law governing all data processing activities.

Leaving the GDPR's trigger conditions to evolve primarily through interpretative practice does not preserve stability. It creates an *interpretative vacuum* in which the practical scope of EU data protection law expands through enforcement dynamics rather than legislative design.

The same structural logic applies to Article 9. The distinction between ordinary personal data and special categories of data performs a comparable boundary-setting function within the Regulation. If inference-based interpretations progressively expand the concept of sensitive data without legislative clarification, the Regulation's internal architecture risks losing the proportional structure originally built into it. Clarifying these trigger conditions, therefore, strengthens the coherence of the GDPR while preserving its foundational protections.

4.1 Actor-Relative Personal Data After *EDPS v SRB* and the Constitutional Stakes of the GDPR Trigger Logic

The Court of Justice's judgment in *EDPS v SRB*[78] does not merely "add colour" to the GDPR's definition of personal data; it sharpens the Regulation's trigger logic by making the practical meaning of identifiability expressly actor-relative and context-dependent. In substance, the Court approaches the threshold question of whether information is personal data by asking who holds the information, what additional means are realistically available to that actor, and what technical, organisational, and legal constraints govern that availability. Two points matter most.

First, pseudonymised information may remain personal data for one actor, typically the controller who holds or can access the additional information, but not constitute personal data for another actor if effective measures prevent that actor from re-identifying the data subject, including by combining it with other datasets. That is a faithful operationalisation of Recital 26 GDPR.[79] and is not a novel invention: it consolidates the trajectory running from *Breyer*[80] to *IAB Europe*[81] and *OC v Commission*[82].

Second, the Court blocks a transparency escape route. For the controller's duty to inform data subjects about recipients[83], identifiability is assessed at the point of collection and from the controller's perspective, irrespective of whether the recipient could identify the data subject after pseudonymisation. Actor relative identifiability, therefore, cannot be used to conceal or downplay disclosure.

These doctrinal moves matter because Article 4 GDPR is architecture, not ornament. The trigger logic serves a constitutional and regulatory purpose. It preserves protection where processing poses real risks to rights and freedoms, while ensuring that the intensity of the GDPR's obligations remains proportionate to the actor's position, the data they hold, and the means reasonably likely to be used. Read that way, actor-relative identifiability does not weaken protection. It helps preserve a risk-based structure in which the Regulation applies where rights are genuinely at stake, without treating every downstream recipient of information that is constrained or effectively non-identifiable as though they stood in the same position as the original controller.

The Digital Omnibus makes that choice unavoidable. It proposes to insert an express actor-relative clarification into Article 4(1) GDPR, phrased negatively: information "is not necessarily personal data for every other person or entity".[84] In Joint Opinion 2/2026, the EDPB and EDPS urge the co-legislators not to adopt that definitional change, arguing that it would narrow the GDPR's scope, enable circumvention, and create inconsistencies across the wider EU legal framework that relies on the concept of personal data.[85]

That position does not merely defend continuity. It also preserves the *status quo* described in the previous section, in which expansive and divergent understandings of identifiability increasingly emerge through guidance, enforcement practice, and institutional drift rather than through legislation. The practical costs are substantial. Businesses face uncertainty about whether the GDPR applies at all, whether pseudonymisation meaningfully changes their legal position, and whether different supervisory authorities will treat the same data architecture

78 Case C-413/23 P, *European Data Protection Supervisor v Single Resolution Board*, EU:C:2025:645.
79 Regulation (EU) 2018/1725, recital 16.
80 *Case C-582/14 Patrick Breyer v Bundesrepublik Deutschland EU*:C:2016:779
81 Case C-604/22, *IAB Europe v Gegevensbeschermingsautoriteit*, EU:C:2024:214.
82 *Case C-479/22 P OC v European Commission EU*:C:2024:215, paras 43–45.
83 Regulation (EU) 2018/1725, art 15(1)(d); by analogy, GDPR art 13(1)(e).
84 European Commission, *Proposal for a Regulation amending Regulation (EU) 2016/679 (Digital Omnibus)* COM(2025) 837 final, 54–55.
85 European Data Protection Board and European Data Protection Supervisor, Joint Opinion 2/2026 on the Proposal for a Regulation as Regards the Simplification of the Digital Legislative Framework (Digital Omnibus), 10 February 2026, paras 10–15.

differently.[86] That uncertainty raises start-up[87] and compliance costs[88], discourages data sharing and privacy-preserving design[89], and distorts innovation incentives by making legal exposure turn less on text and principle than on supervisory attitude.[90] A framework meant to be *risk-based* starts to resemble one in which scope itself expands unevenly and unpredictably.[91]

Identifiability is actor-relative. Whether information qualifies as personal data depends on who holds it, what additional means are realistically available to that actor, and what technical, organisational, and legal constraints govern access to those means.

Pseudonymised data do not have a single legal status for all actors. The same information may remain personal data for the original controller while falling outside the GDPR's practical scope for a downstream recipient that cannot realistically re-identify the data subject.

The Court preserves transparency obligations. A controller cannot rely on downstream non-identifiability to avoid informing data subjects about recipients.

Article 4 performs a threshold function. It is not merely definitional. It structures when the GDPR applies and how proportionately its obligations attach across different actors and data environments.
The real policy dispute is about legal certainty and regulatory breadth. The Omnibus would codify an actor-relative approach, while the EDPB and EDPS prefer to preserve the current framework, despite the uncertainty and divergence it has generated in practice.

The difficulty with the EDPS/EDPB approach is not simply that it resists amendment. It is that, in the name of preventing loopholes, it entrenches a model in which the GDPR's trigger conditions continue to drift toward breadth without clear limiting principles. That is constitutionally and practically problematic. If identifiability is always interpreted at the outer edge of possibility, regardless of the actor's actual position and realistic means, the threshold function of Article 4 begins to collapse. The concept of personal data no longer distinguishes between *materially different risk positions*. It becomes a one-way ratchet.

That approach has consequences. It weakens legal certainty, invites inconsistent enforcement, and makes it harder for organisations to predict when obligations arise and how far they extend.[92] It also places growing pressure on innovation and routine data use, not because the legislature explicitly chose that outcome, but

[86] American Chamber of Commerce to the European Union, 'Advancing EU data and cybersecurity rules through the Digital Omnibus' (Position Paper, 11 March 2026) available at https://www.amchameu.eu/position-papers/advancing-eu-data-and-cybersecurity-rules-through-the-digital-omnibus accessed 2 April 2026; American Chamber of Commerce to the European Union, 'Reducing complexity in Europe's digital rulebook' (24 March 2026) available at https://www.amchameu.eu/news/reducing-complexity-in-europe-s-digital-rulebook accessed 2 April 2026.

[87] SMEunited, 'A Digital Package for Simpler Rules' (Position Paper, 2 February 2026) 1, available at https://www.smeunited.eu/admin/storage/smeunited/smeunited-digitpackage.pdf (accessed 2 April 2026); Allied For Startups, 'A simplified digital single market for innovation and growth: Simplification, digital package and omnibus, Call for Evidence' (14 October 2025) 1–2, available at https://alliedforstartups.org/wp-content/uploads/AFS-Position-Digital-Package-Omnibus.pdf (accessed 2 April 2026).https://www.smeunited.eu/admin/storage/smeunited/smeunited-digitpackage.pdfhttps://alliedforstartups.org/wp-content/uploads/AFS-Position-Digital-Package-Omnibus.pdf

[88] BusinessEurope, 'Making the General Data Protection Regulation More Effective' (Position Paper, February 2026) 1–2.

[89] Ecommerce Europe, 'Position Paper on Online Advertising' (13 October 2025) 1, 4–5, available at https://ecommerce-europe.eu/wp-content/uploads/2025/10/ECOM-position-paper-on-Online-Advertising.pdf (accessed 2 April 2026).https://ecommerce-europe.eu/wp-content/uploads/2025/10/ECOM-position-paper-on-Online-Advertising.pdf

[90] CCIA Europe, 'CCIA Europe Response to the European Commission's Public Consultation on the Upcoming Digital Omnibus Simplification Package: Simplifying the EU digital rulebook for innovation' (October 2025) 1–2, available at https://ccianet.org/wp-content/uploads/2025/10/CCIA-Europe-Response-Digital-Package-Consultation.pdf (accessed 2 April 2026).https://ccianet.org/wp-content/uploads/2025/10/CCIA-Europe-Response-Digital-Package-Consultation.pdf

[91] DIGITALEUROPE, 'Digital simplification package: Our data recommendations' (Position Paper, 4 June 2025) 9–10, available at https://cdn.digitaleurope.org/uploads/2025/06/Digital-simplification-package-Data.pdf (accessed 2 April 2026).https://cdn.digitaleurope.org/uploads/2025/06/Digital-simplification-package-Data.pdf

92 DIGITALEUROPE, 'Digital simplification package: Our data recommendations' (Position Paper, 4 June 2025) 9–10 <https://cdn.digitaleurope.org/uploads/2025/06/Digital-simplification-package-Data.pdf> accessed 2 April 2026.

because soft law and enforcement practices have gradually expanded the regime's perimeter.[93] The result is not stronger data protection in any coherent sense. It produces a system in which the scope of the GDPR becomes unstable across actors and contexts, and in which organisations struggle to calibrate compliance obligations predictably and proportionately. Maintaining that status quo is therefore not neutral. It preserves the very pathologies of breadth, divergence, and uncertainty that now justify legislative clarification in the first place.

4.2 Actor-Relative Identifiability in *EDPS v SRB*

The factual set-up matters because the Court uses it to articulate the doctrinal test. In the SRB "right to be heard" procedure, the Single Resolution Board transmitted comments to Deloitte using an alphanumeric code; only the SRB could link those coded comments back to registrant identity data, and Deloitte did not have access to the relevant database.[94] The Court's key move is to treat pseudonymisation as a contextual barrier whose legal effect depends on whether it actually prevents attribution for the relevant actor.

For the controller (here, SRB) that retains "additional information", pseudonymised comments remain personal data because the controller can still attribute them to a data subject. For the recipient (here, Deloitte), the comments may be non-personal data if technical/organisational measures ensure that Deloitte cannot lift the pseudonymisation or re-identify the person via "other means … such as cross-checking with other factors", so that the person is not (or is no longer) identifiable for that company.[95]

This claim follows the structure already embedded in Recital 16 of Regulation 2018/1725[96] (which mirrors Recital 26 GDPR). Identifiability turns on "all the means reasonably likely to be used", assessed using objective factors (cost/time/technology), and the principles should not apply to information rendered anonymous such that the data subject is not or no longer identifiable. Crucially, the Court rejects an absolutist reading: pseudonymised data are not, in all cases and for every person, personal data. Pseudonymisation can, depending on circumstances, prevent identification for actors other than the controller.[97]

The *SRB* judgment is often summarised as "data can be personal for A but not for B". That is true, but incomplete. The Court emphasises that identifiability analysis is not a free-floating abstraction: it depends on "the circumstances of the processing … in each case". And when assessing whether means are "reasonably likely", the Court accepts that a risk of identification is not "reasonably likely" where the risk is insignificant, including because identification is prohibited by law or impossible in practice due to disproportionate effort in time/cost/labour.

Pseudonymised data are not, in all cases and for every person, personal data. Whether they remain identifiable depends on the actor, the means available to that actor, and the circumstances of the processing.

But the Court simultaneously insists on a downstream sensitivity: where data are put at the disposal of others with means reasonably likely to identify them, "otherwise impersonal" data can become personal for those recipients and (indirectly) for the entity making them available, an idea the Court links to prior case law. This is exactly the point the EDPB/EDPS later used to criticise the Omnibus wording that tries to deny "becoming personal" merely because downstream recipients can identify.

---

[93] EuroCommerce, 'EuroCommerce response to the Call for Evidence on the proposed Communication on Better Regulation' (4 February 2026) 3, available at https://www.eurocommerce.eu/app/uploads/2026/02/20260204-eurocommerce-position-call-for-evidence-better-regulation-final.pdf (accessed 2 April 2026).https://www.eurocommerce.eu/app/uploads/2026/02/20260204-eurocommerce-position-call-for-evidence-better-regulation-final.pdf

94 Case C-413/23 P, *European Data Protection Supervisor v Single Resolution Board*, EU:C:2025:645.

95 *EDPS v SRB*, para 75.

[96] Regulation (EU) 2018/1725 of the European Parliament and of the Council of 23 October 2018 on the protection of natural persons with regard to the processing of personal data by the Union institutions, bodies, offices and agencies and on the free movement of such data, and repealing Regulation (EC) No 45/2001 and Decision No 1247/2002/EC (Text with EEA relevance.)

97 *EDPS v SRB*, paras 73–74.

4.3 The Trigger Logic Consequence: Duties Can Attach Even If the Recipient Cannot Identify

One of the most overlooked parts of *SRB* is its holding on the moment and perspective for applying a controller's information duty to recipients. For Article 15(1)(d) Regulation 2018/1725 (recipients/categories of recipients), the Court holds:

- The duty applies at the time data are obtained, and
- Identifiability must be assessed at collection and from the controller's point of view, before any transfer.

Accordingly, SRB's duty to inform applied irrespective of whether the transferred comments were, from Deloitte's perspective, personal data after pseudonymisation. That is the pivot that keeps the GDPR's trigger logic coherent: actor-relative identifiability calibrates scope, but it does not let controllers "launder" disclosures by choosing recipients who cannot, or contractually promise not to, re-identify.

Figure 4.1. Actor-Relative Identifiability Test

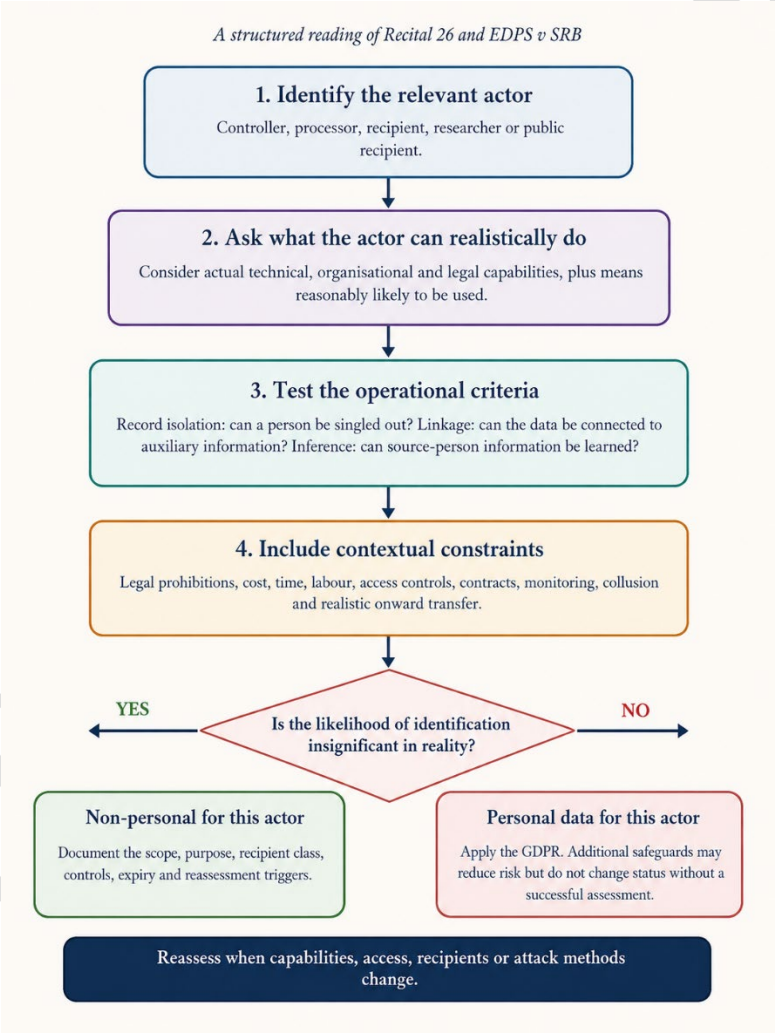

Note: A structured reading of Recital 26 GDPR / Recital 16 Regulation 2018/1725 as operationalised in EDPS v SRB: "Personal data" is assessed per actor and per context, but disclosure risk and realistic downstream identification remain legally salient.

Figure 4.2. The Same Pseudonymised Dataset Across Actors

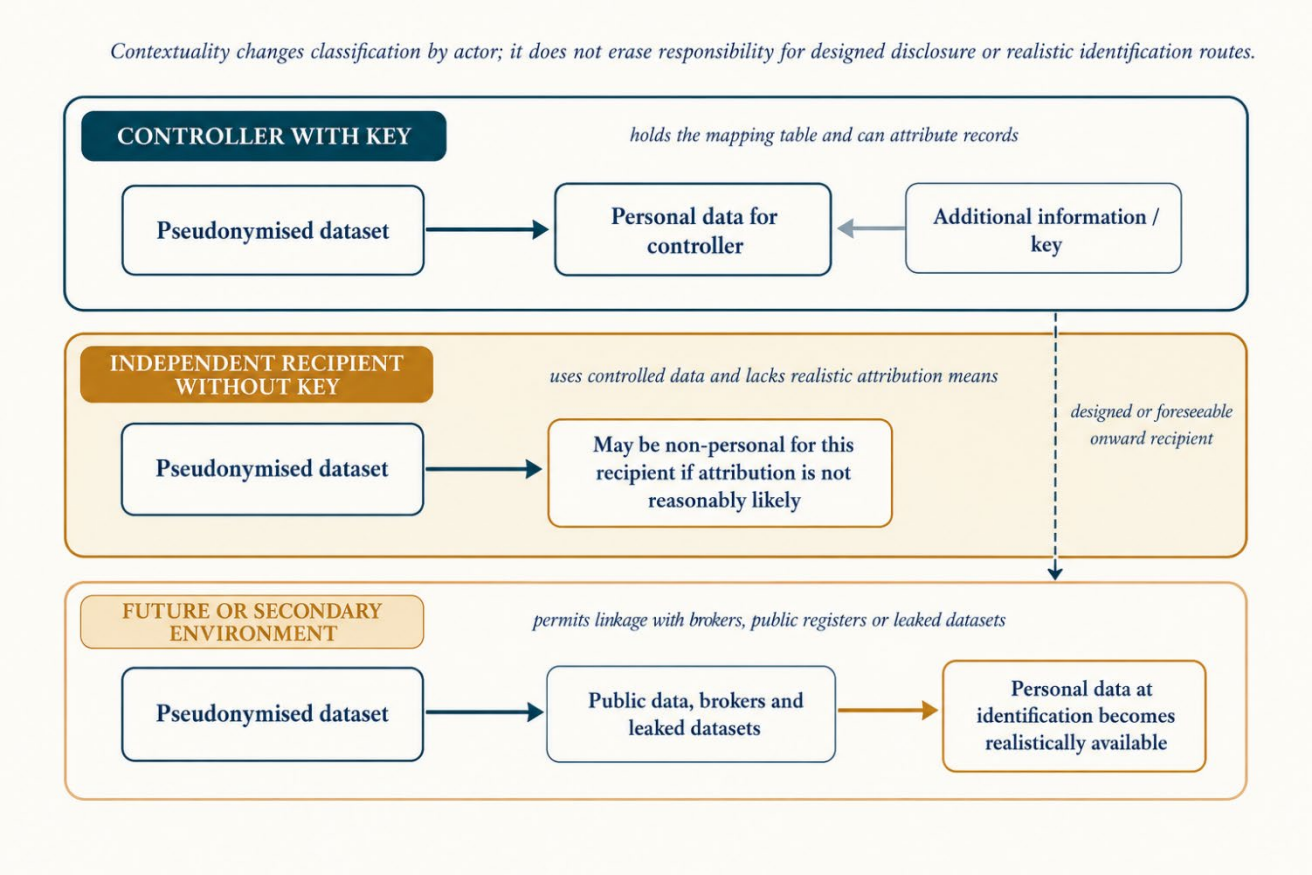

Notes: SRB-style actor relativity: pseudonymised information can be personal for the key-holder and (depending on context) non-personal for a constrained recipient; uncontrolled or enriched contexts can reactivate identifiability.

4.4 Guidelines 02/2026: Contextuality Conceded, Certainty Withheld

Guidelines 02/2026 remain a consultation text rather than final guidance.[98] Their institutional status nevertheless matters because the EDPB adopted them after the Board and Supervisor had urged the co-legislators to leave anonymisation to forthcoming guidance.[99] The text therefore provides the best available account of the Board's post-SRB position.

The central concession is explicit. Anonymity may vary between entities; each relevant entity's perspective must be assessed; and the contextual approach reflects the full legal standard. The simplified approach may still serve a precautionary or administrative purpose, but the Guidelines acknowledge that it can treat genuinely anonymous information as personal in a particular context.[100]

The concession is not the same as an operational safe route. Paragraph 21 asks whether an identifying downstream recipient 'cannot be ruled out'; paragraph 30 brings cybercriminals, rogue employees and foreign intelligence services into the universe of relevant actors; and the simplified approach assumes the broadest status whenever contextual analysis becomes difficult. Those moves can recreate an absolute test through the assessment methodology after the Guidelines have accepted contextuality as the legal principle.[101]

The operational framework uses No Record Isolation, No Linkage and No Inference, but those criteria do not deliver a decision rule. Even where an actor satisfies all three, paragraph 52 says that the information 'may' be

[98] EDPB, Public Consultation on Guidelines 02/2026 on Anonymisation, consultation period 8 July to 30 October 2026.

[99] European Data Protection Board and European Data Protection Supervisor, Joint Opinion 2/2026 on the Proposal for a Regulation as regards the simplification of the digital legislative framework, adopted 10 February 2026, especially paras 13–25 and 96–117.

[100] EDPB, Guidelines 02/2026, which distinguish a contextual approach reflecting the full legal standard from a simplified approach that may treat information as personal even where it is anonymous in the relevant context.

[101] 'Problems with EDPB Guidelines 02/2026 on Anonymisation: Critical Analysis for Consultation Response' (working draft, July 2026, on file with the author), sections I-III and IX; EDPB, Guidelines 02/2026, especially paras 21, 30, 48 and 52.

regarded as anonymous. Failure of any criterion also does not establish that the data are personal; it triggers further analysis. The framework therefore leaves both routes open and places the decisive threshold in an undefined residual judgment. That may suit supervisory caution, but it does not give organisations a stable legal conclusion.[102]

The Guidelines also risk conflating the Article 4 scope inquiry with an Article 32 threat model. Illegal intrusion, a rogue employee, a hypothetical leak and future technical advances all matter to security and lifecycle review. They should affect identifiability only where evidence makes the route reasonably foreseeable for the relevant actor. Otherwise, every useful dataset remains perpetually personal because some unlawful or future actor might identify somebody.

Two further examples show why legislation needs a clearer threshold. The Guidelines sometimes treat uniqueness or singling out within a visual or statistical field as sufficient without demonstrating how the actor can connect the distinction to a natural person in a way that permits meaningful treatment. They also suggest that a dataset is anonymous only if the technique succeeds for every included individual. A single unresolved outlier should remain protected and quarantined; it should not automatically determine the status of every separable record. The legal test should ask what can realistically be attributed, by whom and in which operational unit.[103]

The anti-circumvention rules answer the strongest loophole concerns. A processor acting for a controller does not receive an independent anonymity perspective merely because it lacks the key. An independent controller, including a research institute that determines its own purposes and means, may require a separate assessment. The distinction follows function and realistic capability, not labels.[104]

Onward transfer remains part of the analysis. Where an entity can reasonably disclose the information to a third party that possesses identification means, the transfer may render the data personal for the recipient. It may also affect the transferring entity's assessment. Actor, purpose, access mode and circulation therefore bound contextual anonymity. It is not a permanent status that survives uncontrolled release.

The online-advertising example is equally important. Pseudonyms derived from device fingerprints and used to link behaviour across services, build profiles and select advertisements identify users for GDPR purposes. The Guidelines therefore reject the proposition that contextuality turns ordinary behavioural tracking into anonymous processing.[105]

Contracts occupy a middle position. A contractual prohibition cannot, by itself, transform identifiable data into anonymous data. Reliable, verifiable and enforceable restrictions can nevertheless reduce the means reasonably likely to be used when they operate alongside access controls, logging, separation of duties, monitoring and sanctions. The legal question concerns the architecture as a whole, not a binary choice between 'contract' and 'technology'.[106]

The No Inference analysis also helps separate general learning from source-person leakage. Population-level knowledge does not become personal data merely because an actor later applies it to a new person. A model, coefficient, aggregate, or synthetic dataset requires testing to see whether queries can reveal specific and

[102] EDPB, Guidelines 02/2026, sections on No Record Isolation, No Linkage and No Inference, together with the discussion of realistic means, auxiliary information, contracts, onward transfer and periodic reassessment.

[103] 'Problems with EDPB Guidelines 02/2026 on Anonymisation' (working draft, July 2026, on file with the author), sections IV-VII, addressing singling out, mixed datasets, aggregate information and AI models; EDPB, Guidelines 02/2026, especially paras 24, 36, 52 and 82.

[104] EDPB, Guidelines 02/2026, examples distinguishing a processor acting on a controller's behalf from an independent controller such as a research institute. The processor does not obtain a separate anonymity perspective merely because it lacks the key.

[105] EDPB, Guidelines 02/2026, online-advertising example concerning pseudonyms derived from device fingerprints and used to link behaviour across services, profile users and select advertisements.

[106] EDPB, Guidelines 02/2026, discussion of legal and contractual constraints. A contract does not by itself create anonymity, but reliable, verifiable and enforceable restrictions may affect which means are reasonably likely when combined with effective technical and organisational controls.

meaningful information about a person represented in the source data. Personal-data lineage triggers scrutiny; it does not create irreversible legal taint.[107]

These propositions substantially weaken the civil-society claim that an entity-relative approach necessarily creates a tracking or pseudonymisation loophole. The restrictive and internally open-ended parts of the Guidelines do not reverse that concession. They instead show why the legislature must translate contextuality into an operative threshold, objective factors, anti-circumvention rules, and a reliance mechanism, rather than leave the result to a methodology that can recreate absolutism through precaution.

4.5 Evidence from Real Architectures: Operationalising SRB without Collapsing Protection

The joint SAMMAN and Team Blaeu report provides the kind of evidence that a workable legal test should rely on. It applies a two-step framework to five company-provided architectures. Step one asks the SRB question from the recipient's perspective: can the recipient lift the measures or identify them by cross-checking with other factors? Step two adds the stricter 2014 anonymisation-opinion tests on singling out, linkability and inference as a precautionary overlay. The framework is deliberately designed to return a negative answer.[108]

The results matter because they do not all point toward deregulation. Four architectures are assessed as most likely to be anonymous for a defined recipient and on defined facts. At the same time, the source data remains personal to the actor who retains the means of identification. One architecture fails. The counter-example demonstrates that pseudonymisation should have legal significance without being treated as a magic label. The legal question is what the combined architecture actually prevents, for whom and for how long.

Table 4.1. Five Real Architectures Under a Recipient-Specific Assessment

Use case	Converging safeguards	Reported recipient conclusion	Legal design lesson
Scrubbed prompts to a generative AI provider	Identifier removal; technically enforced zero retention; no logging, training or reuse; encrypted transfer.	Most likely anonymous for the AI provider; personal for the legal intelligence provider. Attachments require a separate assessment.	The contract alone is insufficient; technical enforcement and the absence of a retained matching dataset are decisive.
Live traffic maps from vehicle signals	Random batching; timing gaps; deletion of backend copies; conversion of GPS points to road segments; swarm aggregation.	Most likely anonymous to the third-party provider once the minimum contributor threshold is met.	Spatial and temporal transformation must be assessed together; sparse routes require contextual thresholds.
Cloud anomaly detection	Salted hashing and environmental separation, but the controller retains the key, and the recipient remains close to the means.	Most likely not anonymous under the framework.	Architectural proximity and surviving reversibility can defeat otherwise strong privacy-enhancing techniques.
Aggregated audience statistics	Aggregation; panel roster, mapping and hybridisation logic withheld; contractual restrictions support technical separation.	Most likely anonymous for the publishing platform.	Withheld mapping and aggregation carry the result; contractual limits reinforce but do not independently establish anonymity.
Synthetic data for ad prediction	Pseudonymisation; group counts; minimum cell sizes; differential privacy; tracked privacy budget; synthetic reconstruction.	Most likely anonymous for the organisation using the synthetic set; source records remain personal.	A forward-only transformation and disciplined privacy budget can sever record-level linkage, subject to continuing statistical controls.

[107] EDPB, Guidelines 02/2026, No Inference section, distinguishing population-level learning from a specific and meaningful inference connected to a natural person represented in the source data.

[108] Rob van Eijk and Benjamin de Vanssay, From Pseudonymised to Anonymous Data: A practical assessment framework under the SRB ruling, applied to real use cases provided by European companies (SAMMAN and Team Blaeu, June 2026), sections 2–7.

Across the cases, seven recurrent design principles emerge. The law should reward architectural separation; recognise the convergence of multiple barriers; use the organisation's actual technical, legal and organisational capabilities as the reference point; treat contractual and commercial constraints as reinforcing rather than standalone safeguards; require factual mitigation of context-specific risks; insist on recipient-specific assessment; and recognise that layered protection is stronger than any single technique.

Those principles also identify the limits of an assurance regime. The report rests on five use cases supplied by interested companies, and four pass the authors' framework. The conclusions should therefore be treated as working hypotheses, not categorical proof. Independent attack testing, supervisory scrutiny and a certification pilot are appropriate ways to confirm or correct them. That caution strengthens rather than weakens the legislative case: the Union needs a common process for converting evidence into reliable legal assurance.[109]

The use cases expose the central weakness of the proposed Article 29a. A non-binding opinion can describe good practice, but it cannot tell a recipient that an architecture accepted in one Member State will receive the same legal treatment in another, or that reasonable reliance will be respected if the state of the art later changes. The better model is legislative criteria plus an EDPB-coordinated assurance system, with ongoing duties to monitor changes in capability, onward transfers, and new linkage routes.

109 Rob van Eijk and Benjamin de Vanssay, From Pseudonymised to Anonymous Data: A Practical Assessment Framework Under the SRB Ruling, Applied to Real Use Cases Provided by European Companies (SAMMAN and Team Blaeu, June 2026), section 8.

Section 5: Article 9, Inference, and the Omnibus Recalibration of the GDPR's Internal Boundary

If Article 4(1) defines the GDPR's external boundary, Article 9 performs a different but equally important function within the regime itself. It identifies those forms of processing that merit heightened legal constraints because of the kind of information involved and the seriousness of the risks that may follow from its use. On that reading, Article 9 is not just a prohibition plus exceptions. It is an internal boundary-setting device. It preserves proportionality under the GDPR by distinguishing ordinary personal data from categories of data whose processing may pose especially significant risks to fundamental rights and freedoms. A framework designed to calibrate obligations by reference to risk begins to lose that structure when questions of scope drift outward and are resolved unevenly.[110]

Furthermore, a framework presented as risk-based may, in practice, generate the opposite impression when its scope becomes uncertain and expansive. In that setting, compliance no longer looks like proportionate calibration within a stable legal perimeter, but like adaptation to an increasingly burdensome and unpredictable boundary of application. At that point, the GDPR starts to resemble a regime in which the boundary of application itself expands unpredictably, rather than one in which clearly applicable obligations are modulated proportionately by risk.[111]

That point matters because the present debate is not about whether special category data deserve strong protection. They plainly do. The question is whether Article 9 can continue to perform its internal boundary-setting role if sensitivity is allowed to expand across increasingly indirect, inferred, incidental or probabilistic forms of revelation. Once that expansion goes too far, Article 9 ceases to distinguish between exceptional and ordinary cases. It becomes an undifferentiated super-trigger within the GDPR itself. That does not simply increase protection. It risks misallocating compliance effort.[112]diluting legal certainty and weakening the Regulation's own proportional structure[113].

This internal boundary has come under increasing strain. Recent case law and regulatory practice have expanded the situations in which Article 9 may be engaged, not only where sensitive data are directly collected, but also where special category status is said to arise through inference, combination, probability, or incidental exposure.[114]

That development reflects a legitimate concern. Controllers should not be able to evade Article 9 simply by avoiding explicit collection while deriving the same sensitive conclusions through analytics, profiling, or model-based enrichment. But the symmetrical risk is now hard to ignore. If Article 9 expands without sufficient legislative or interpretative discipline, it begins to lose its internal boundary-setting role. Instead of singling out genuinely high-risk forms of processing, it begins to function as an undifferentiated super-trigger that imposes exceptional restrictions on a widening range of data uses.

---

[110] Claudia Quelle, 'Enhancing Compliance under the General Data Protection Regulation: The Risky Upshot of the Accountability and Risk Based Approach' (2018) 9(3) European Journal of Risk Regulation 502.

111 CNIL, 'AI and GDPR: The CNIL Publishes New Recommendations to Support Responsible Innovation' (guidance, accessed 2 April 2026).

[112] BusinessEurope, 'Making the General Data Protection Regulation More Effective' (Position Paper, February 2026) 1–2, available at https://www.businesseurope.eu/wp-content/uploads/2026/02/2026–02–11-Making-the-General-Data-Protection-Regulation-more-effective.pdf (accessed 2 April 2026).https://www.businesseurope.eu/wp-content/uploads/2026/02/2026-02-11-Making-the-General-Data-Protection-Regulation-more-effective.pdf

[113] Claudia Quelle, 'Enhancing Compliance under the General Data Protection Regulation: The Risky Upshot of the Accountability and Risk Based Approach' (2018) 9(3) European Journal of Risk Regulation 502; Milda Macenaite, 'The "Riskification" of European Data Protection Law through a Two Fold Shift' (2017) 8(3) European Journal of Risk Regulation 506.

114 Case C-184/20, *OT v Vyriausioji tarnybinės etikos komisija*, EU:C:2022:601; Case C-252/21, *Meta Platforms Inc and Others v Bundeskartellamt*, EU:C:2023:537; Case C-21/23, *Lindenapotheke*, EU:C:2024:846; Case C-446/21, *Maximilian Schrems v Meta Platforms Ireland Ltd*, EU:C:2024:834; Case C-65/23, *MK v K GmbH*, EU:C:2024:1051; Case C-492/23, *X v Russmedia Digital SRL and Inform Media Press SRL*, EU:C:2025:935.

That is not a marginal drafting concern. It affects the coherence of the GDPR's architecture. When too many forms of data are treated as sensitive, the law loses its capacity to differentiate between ordinary and exceptional cases. Resources are diverted away from genuinely high-risk uses. Compliance burdens rise even where the connection to actual rights-based harm is remote or speculative. Innovation-by-design becomes harder, not because the law is protecting core interests more effectively, but because its internal thresholds no longer perform a stable sorting function.

5.1 Inference-Based Sensitivity in the Case Law and Guidance

EU law already makes clear that Article 9 cannot be neutralised simply by avoiding the direct collection of manifestly sensitive inputs while engineering the same result downstream through inference, aggregation, or contextual revelation.[115] In *OT v Vyriausioji tarnybinės etikos komisija*, the Court held that the publication of a spouse's name could amount to processing of special category data because, in context, it was liable indirectly to disclose sexual orientation.[116] In *Meta Platforms and Others v Bundeskartellamt*, the Court went further. It held that Article 9 may be engaged where online conduct, app or website visits, and the use of embedded tools reveal protected characteristics through the combination of behavioural signals and contextual data. It stressed that the prohibition in Article 9(1) applies regardless of whether the sensitive conclusion was the controller's aim or whether the conclusion is accurate.[117] *Lindenapotheke* confirms the same trajectory in the health-data setting: data entered when ordering pharmacy-only medicinal products online may qualify as "data concerning health" because, through collation or deduction, they are capable of revealing information about the data subject's health status, even where the relevant conclusion is probabilistic rather than certain.[118] Most recently, *Schrems v Meta Platforms Ireland reaffirmed the Court's strict approach to Article 9(2)*(e), holding that isolated public statements by a data subject do not collapse Article 9 protection for later, distinct advertising uses, and that "manifestly made public" requires an intentional, clear, and affirmative act directed at making the relevant data public.[119]

The doctrinal point is therefore sharper than a bare statement that Article 9 covers "inferred" data. The Court's case law now treats Article 9 as concerned with substantive revelatory effect, not merely with the formal category of the input data. If ordinary personal data, viewed singly or in combination, discloses information about health, sexuality, religion, political opinion, or other protected matters, Article 9 may be triggered even where the controller did not explicitly solicit that information and even where the resulting conclusion is probabilistic, indirect, or generated only through contextual processing.[120] The practical consequence is that the Article 9 boundary is increasingly shaped by the inferential environment of the processing operation rather than by the face value of the collected data alone.

The same logic appears in supervisory guidance. The EDPB's Guidelines 8/2020 on targeting social media users state expressly that assumptions or inferences regarding special category data may themselves constitute special category data.[121] Political preference inferred from page visits is the Board's own example. That position is entirely understandable in targeting and profiling environments, because those systems are often designed precisely to infer latent traits from observable conduct. But it also shows why Article 9 is now under pressure. Once inferential sensitivity is recognised, the law needs limiting principles if Article 9 is to remain a meaningful internal boundary rather than an expanding category without stable edges.

---

[115] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) [2016] OJ L119/1 art 9.

[116] *OT v Vyriausioji tarnybinės etikos komisija* (*Case C-184/20*) EU:C:2022:601, Court of Justice of the European Union, 1 August 2022.

[117] *Meta Platforms Inc, Meta Platforms Ireland Ltd and Facebook Deutschland GmbH v Bundeskartellamt (Case C-252/21*) EU:C:2023:537, Court of Justice of the European Union, 4 July 2023.

[118] *ND v DR (Lindenapotheke)* (C-21/23) EU:C:2024:846, paras 83–90.

[119] *Schrems v Meta Platforms Ireland Ltd (C-446/21*) EU:C:2024:834, especially paras 74–80 on Article 9(2)(e).

[120] GDPR art 9(1); *Meta Platforms v Bundeskartellamt (C-252/21*) EU:C:2023:537, paras 68–70; *Lindenapotheke* (C-21/23) EU:C:2024:846, paras 83–90.

[121] European Data Protection Board, Guidelines 8/2020 on the targeting of social media users (Version 2.0, adopted 13 April 2021) 32–34.

Lindenapotheke shows how far this logic can travel. There, the Court held that data entered when ordering pharmacy-only medicinal products online could constitute data concerning health, because such data may reveal information about the health status of the person concerned through inference or deduction. The Court's approach, therefore, reaches beyond explicit disclosure and into the realm of *reasonably derivable sensitivity*. Yet the case law does not support an unlimited collapse of boundaries. In *Schrems (Communication of data to the general public)*, the Court did not accept that a public statement by the data subject about his sexual orientation gave Meta a general licence to process other sexuality related data for personalised advertising. Article 9 is therefore already stringent, but it is not meant to become a rule under which any remotely revealing datum is automatically sensitive in every setting and for every purpose.[122]

Read together with Recital 51 of the GDPR, the *Vyriausioji tarnybinės etikos komisija* confirms that Article 9 is not a technical add-on but a regime of enhanced protection for particularly sensitive data; in other words, a protective intensifier designed to process forms that pose *qualitatively* different risks. This regime is justified because such processing may pose significant risks and constitute a particularly serious interference with privacy and data protection rights.

At the same time, that rationale implies a *limiting* principle. The stronger the legal consequences attached to Article 9 classification, the greater the need to ensure that the category remains tethered to a meaningful conception of *sensitivity* rather than expanding through increasingly attenuated inferential chains. The result is an important but difficult position. Article 9 is already inference-sensitive. That much is settled. But inference sensitivity cannot mean that any datum becomes sensitive whenever it might, in some chain of reasoning, contribute to an eventual sensitive attribution. If that were the rule, Article 9 would cease to distinguish exceptional cases from ordinary processing. It would become structurally *unstable*.

5.2 Why the Omnibus Approach Is Necessary

The Omnibus debate should be understood as a stress test of Article 9's internal boundary, not as a retreat from the principle that genuinely sensitive data merit enhanced protection. The GDPR itself explains why special categories receive that treatment. Recital 51 states that data particularly sensitive in relation to fundamental rights and freedoms merit specific protection because the context of their processing may create significant risks, while Recital 52 permits derogations only in limited situations.[123] The Court has framed the point in constitutional terms: in *OT*, it stated that Article 9(1) seeks to ensure enhanced protection against processing liable, by reason of the particular sensitivity of the data concerned, to constitute a particularly serious interference with Articles 7 and 8 of the Charter.[124] The real issue, therefore, is not whether special category data deserve strong protection. They do. The issue is whether the present doctrinal trajectory still preserves a workable distinction between processing that is genuinely special in its object, purpose, or practical focus and the now-routine operation of digital systems in which some degree of sensitive meaning may emerge only residually, contextually, or inferentially.

That distinction matters because the GDPR is not built as an undifferentiated precautionary code. Its architecture is risk-based. Recitals 74 to 76 require controllers to calibrate compliance measures with reference to the nature, scope, context, and purposes of processing, and to risks of differing likelihood and severity.[125] Yet the recent case law on Article 9 pushes in the opposite direction at the threshold stage. *OT* treats relationship data as capable of disclosing sexuality. *Meta v Bundeskartellamt* confirms that behavioural and contextual data can become special category data through combination, and that Article 9(1) operates independently of both the controller's aim and the correctness of the sensitive conclusion.[126] *Lindenapotheke* extends the same logic to health data, holding that

[122] Court of Justice of the European Union, Press Release No 159/24, Judgment in Case C-21/23 *Lindenapotheke*, 4 October 2024.
[123] GDPR recitals 51–52.
[124] *OT v Vyriausioji tarnybinės etikos komisija (C-184/20)* EU:C:2022:601, para 100 and related reasoning on the seriousness of the interference.
[125] GDPR recitals 74–76.
[126] *Meta Platforms v Bundeskartellamt (C-252/21)* EU:C:2023:537, paras 68–70, 89.

an intellectual operation of collation or deduction can suffice, even without a diagnosis or prescription.[127] *Schrems v Meta Platforms Ireland* then reads Article 9(2)(e) narrowly, insisting that occasional public disclosure by the data subject does not dissolve the protective structure for later processing.[128] Read together, those authorities point toward an Article 9 threshold that increasingly captures data because sensitive meaning can be inferred, assembled, probabilistically attributed, or contextually exposed through the processing environment.

Figure 5.1. Article 9 Trigger Logic

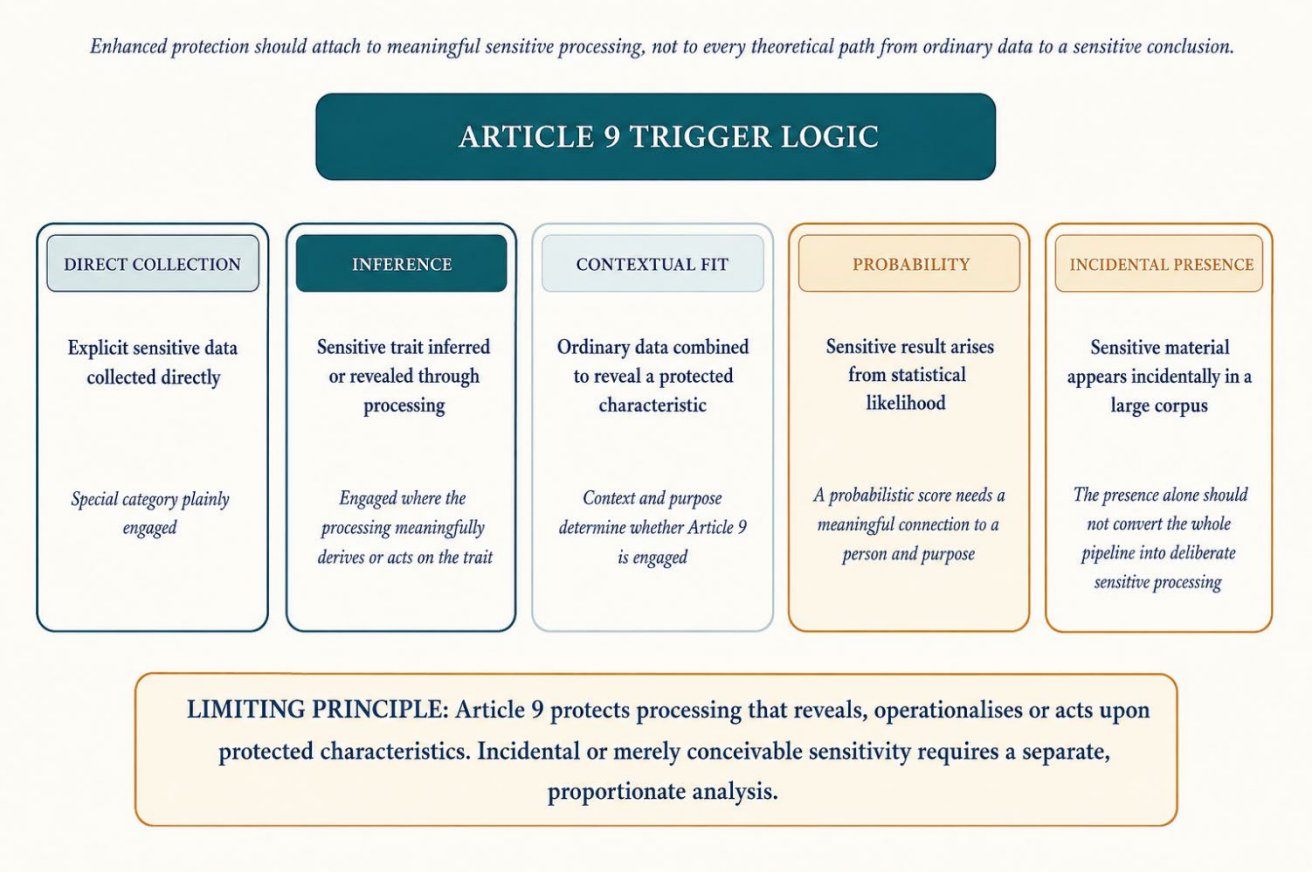

Accordingly, Article 9(1) ***cannot be treated as beyond reform***. If the threshold of sensitivity expands without firmer limiting principles, the pressure is merely displaced onto Articles 9(2) and 9(5), which were designed as exceptions and safeguard mechanisms, not as shock absorbers for an overextended trigger. The problem becomes acute in AI and large-scale data environments. The EDPB's 2024 Opinion on AI models recognises that AI development and deployment raise systemic, abstract, and novel data-protection issues, and that supervisory authorities must assess the nature of the data, the context of the processing, and the possible further consequences for data subjects.[129] Recent French guidance points in a more operational direction. The CNIL states that if, despite appropriate measures, a controller incidentally and residually processes sensitive data that it did not seek to collect, that fact alone should not be treated as unlawful; if the controller later becomes aware of that sensitive processing, it should erase it as far as possible.[130] That does not abandon Article 9. It reflects an attempt to distinguish targeted exploitation of sensitive data from incidental or non-targeted encounters with sensitive material in complex technical systems.

The Commission's Omnibus approach is best understood not as deregulation, but as an effort to restore doctrinal manageability to a provision whose current trajectory risks turning an internal intensifier into a pervasive super-

[127] *Lindenapotheke* (C-21/23) EU:C:2024:846, paras 83–90.
[128] *Schrems v Meta Platforms Ireland Ltd (C-446/21)* EU:C:2024:834, paras 74–80.
[129] European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (18 December 2024) paras 7–8, 60–67, 133–135.
[130] CNIL, The legal basis of legitimate interest: focus sheet on the measures to implement in the case of data collection by web scraping (5 January 2026).

trigger. A regime that treats every residual, probabilistic, or incidental encounter with potentially sensitive information as functionally equivalent to deliberate sensitive profiling does not preserve proportionality. It risks converting Article 9 from a mechanism of differentiation into a blunt prohibition that ordinary digital development, data sharing, and even privacy-preserving system design must constantly struggle to escape. The better reading is not that Article 9 should retreat from genuinely harmful sensitive processing, but that its threshold should remain capable of distinguishing real, rights-significant sensitive processing from the background inferential noise of contemporary data systems. That is the legal and regulatory space into which the Omnibus intervention properly fits.

5.3 Article 9(2)(k), Paragraph 5, and a Workable Compliance Path for AI

The proposed new derogation in Article 9(2)(k) is therefore best understood as a constructive and necessary response. It permits processing in the context of developing and operating an AI system or model, subject to the conditions set out in paragraph 5. Structurally, that matters because it acknowledges a practical truth that the current framework handles poorly: AI pipelines may contain residual special category data even where that data are not being sought for their own sake and are not necessary in any strong purposive sense. The proposal creates a lawful route for dealing with that reality without normalising unrestricted sensitive processing.

That is precisely the right direction. A well-designed derogation of this kind does not trivialise sensitivity. It subjects processing to conditions, duties of avoidance, removal where feasible, and protective obligations where removal would require disproportionate effort. In that respect, it preserves the normative seriousness of special category data while refusing to treat all incidental encounters with such data as equivalent to deliberate, high-salience exploitation.

The criticisms made in Joint Opinion 2/2026 do not undermine that core insight. The EDPB and EDPS are right to insist that the legislative text should speak clearly and that incidental or residual processing should be expressly reflected in the enacting terms. They are also right to resist any reading of "operation" that would turn deployment into a broad permission to solicit or harvest special category data through prompts or interface design. But those are arguments about delimitation, not about whether Article 9(2)(k) should exist. Properly understood, they support a more precise Omnibus text rather than the abandonment of the Omnibus approach.

This approach also applies to the accompanying amendment to paragraph 5. Its purpose should be to translate this recalibration into proportionate, workable operational safeguards. The law should require controllers to avoid collecting special category data where possible, to remove it when identified and where reasonably practicable, and to protect it where immediate removal would be disproportionate or technically disruptive. That is a recognisable form of risk governance. It focuses on what controllers should do when sensitive data appears, rather than pretending that modern AI systems can always guarantee the complete absence of such data.

5.4 Additional Amendments and the Case for Fuller Alignment

On that basis, support for the Omnibus should extend beyond Article 9(2)(k) alone. The proposed Article 9(2)(l), permitting biometric verification where the biometric data or the means necessary for verification remain under the sole control of the data subject, is also sound. It enables secure and user-controlled authentication without endorsing centralised biometric repositories or unnecessary institutional control over biometric identifiers. It reflects an approach to sensitive data protection that is both rights-conscious and technologically realistic.

More broadly, there is force in extending Article 9(2) to include additional derogations based on legitimate interests and contractual necessity, subject to the safeguards in paragraph 5. The present asymmetry between Articles 6 and 9 is increasingly difficult to defend in contexts where certain services, safety functions, insurance processes, health-related systems, or user-requested features may require strictly necessary processing of special category data even though consent is not the only or most appropriate legal basis in practical terms. A carefully

constrained Article 9(2)(m) and 9(2)(n) would not dissolve Article 9's protections. It would align the Regulation's internal logic more closely with the realities of modern data processing while preserving heightened safeguards where sensitivity genuinely matters.

The same is true of a more workable reformulation of Article 9(2)(g). If the public-interest ground remains so tightly tied to specific legislative prescriptions and detailed prior safeguards that it is rarely usable in practice, the law creates a theoretical derogation without practical effect. Recasting that ground in a way that remains anchored in public interest but does not render its invocation unrealistically burdensome would improve legal certainty and preserve flexibility for legitimate governance purposes.

5.5 Restoring Article 9 to Its Proper Role

Article 9 can remain exceptional only if it continues to serve a genuine boundary-setting function within the GDPR. Just as Article 4 determines when the regime is triggered, Article 9 determines when the law moves from ordinary personal data to a more restrictive and specially protected category. That internal trigger logic matters. It allows the Regulation to preserve stronger protection for processing that creates a qualitatively different level of risk, while maintaining proportionality for the broader universe of personal data processing. If Article 9 loses its sorting function, the Regulation's architecture begins to distort from within.

> **"Article 9 loses its boundary function when it no longer distinguishes incidental sensitivity from the deliberate revelation or use of protected characteristics."**

That is why the present debate should not be framed as a choice between strong rights protection and innovation. The real issue is whether the law still distinguishes in a disciplined way between genuinely sensitive processing and the increasingly common reality of digital systems in which sensitive traits may arise through inference, probability, enrichment, or incidental exposure. To deny that inferential sensitivity matters would understate the risks that Article 9 exists to address. But to allow every residual or probabilistic connection to protected characteristics to pull processing into Article 9 would be to stretch the provision beyond its rationale and to turn an exceptional safeguard into a routine super-trigger. Once that happens, the law no longer concentrates its force where the risks are highest. It begins to flatten important differences of context, function, and purpose.

The case for the Omnibus approach is strongest when framed in those terms. Properly understood, it is not an attempt to hollow out Article 9, but to restore coherence to its internal trigger logic. The aim is not to deny that sensitive inferences can arise without explicit collection. It is to ensure that the legal consequences remain calibrated to what the processing is actually doing, why the data are being used, how realistically protected characteristics are being revealed or operationalised, and what safeguards structure the processing environment. Article 9 should therefore continue to capture genuinely sensitive processing, including cases where systems are in substance revealing, using, or acting upon protected characteristics linked to a reasonably identifiable person. But it should do so through concepts and derogations that remain proportionate, workable, and attentive to context. Otherwise, the law risks treating residual sensitivity, incidental emergence, and deliberate exploitation as legally interchangeable, weakening the very structure that makes heightened protection meaningful in the first place.

5.6 Legislative Choices for the Digital Omnibus Reform

The Omnibus proposal does something rare: it tries to legislate directly on the trigger condition. The proposed addition to Article 4(1) GDPR is:

> *"Information relating to a natural person is not necessarily personal data for every other person or entity…" and "Information shall not be personal for a given entity where that entity cannot identify…" and "Such information does not become personal for that entity merely because a potential subsequent recipient has [means]…"*

As a descriptive statement of SRB's actor-relative logic, the first two sentences are recognisably connected to the Court's approach. But the last sentence ("does not become personal for that entity merely because…") is the most constitutionally loaded: it risks severing the link between downstream identifiability and upstream accountability that SRB itself preserves through its references to prior case law on disclosure to actors with identifying means. The EDPB/EDPS Joint Opinion, therefore, argues that the proposal goes beyond codification and would narrow the concept, produce confusion through "negative definition", and create loopholes. It strongly urges co-legislators not to adopt the definitional change and stresses the knock-on effects, as other EU acts align with or reference the GDPR definition.

The core trade-off is between legislative calibration and interpretative drift. The key is not "change vs no change", but whether the trigger condition is controlled by democratic design or by emergent enforcement dynamics. Against this background, the Omnibus raises a question that is more structural than textual. In a digital economy defined by constant, high-volume, multi-actor data processing, the GDPR cannot depend on unstable trigger conditions and still claim to operate as a coherent risk-based regime. Supervisory resources are finite. Compliance decisions are made at scale and in advance. Accountability therefore depends on thresholds that are sufficiently clear to guide behaviour before the fact, sufficiently principled to concentrate legal force where risks to rights and freedoms are genuinely serious, and sufficiently stable to prevent the regime's scope from being expanded in practice through interpretative drift alone. The real issue is not whether Article 4 may be clarified. It is whether the trigger logic of EU data protection law will remain underdefined and increasingly shaped by emergent enforcement dynamics, or be placed on a sounder legislative footing that better reconciles protection, proportionality, and legal certainty. At this scale, trigger logic is no longer a technical definitional issue. It is part of the regime's governance infrastructure.

If the trigger conditions of the GDPR remain underdefined in legislation, their practical meaning will continue to be shaped through guidance, enforcement and institutional practice. That carries high costs. It weakens legal certainty, entrenches divergent approaches to identifiability, increases compliance burdens, and allows the Regulation's effective boundary to expand without a clear legislative mandate. The case is therefore not against legislative clarification, but for doing it properly. The legislature should stabilise Article 4 in a way that reflects the Court's actor-relative and contextual approach to identifiability while preserving the Regulation's risk-based structure. In practice, that means drafting that: (a) anchors identifiability to Recital 26's language of means reasonably likely to be used, including singling out; (b) preserves downstream disclosure responsibility where making data available creates a realistic identification risk for others; and (c) prevents artificial avoidance through nominal separation of keys, affiliates, processors or contractual arrangements that leave re-identification realistically possible.

Section 6: Specific Issues — Rebutting the Status Quo and the Cost of Inaction

The Omnibus debate is often framed as a choice between strong protection and deregulation. That frame has become harder to sustain after Guidelines 02/2026. The EDPB now accepts that anonymity may vary by entity and that the contextual approach reflects the legal standard. The remaining disagreements concern drafting, burden, anti-circumvention, onward transfer, assurance and institutional control. Those are serious issues, but they are questions about how to legislate contextuality, not whether contextuality is permissible.[131]

Civil-society organisations retain the right to resist tracking, profiling, strategic ignorance and nominal corporate separation.[132] The error lies in treating those risks as reasons to deny recipient-relative classification altogether. The EDPB's own examples show a narrower route: cross-service device-fingerprint profiling remains personal data.[133] A processor inherits the controller's perspective.[134] Realistic disclosure to an identifying recipient matters, and a contract cannot manufacture anonymity without effective controls.[135]

The EDPB's position also changes the institutional argument. In February, the Board and Supervisor asked the legislature to leave the issue to forthcoming guidance. The resulting guidance rejects the absolute approach defended by several opponents of Article 4 reform. Civil society can challenge the Commission's wording and demand stronger safeguards. It can no longer plausibly present contextual anonymity itself as alien to the GDPR.

Across the pressure points below, inaction does not preserve a stable equilibrium. It leaves essential questions to soft law, divergent enforcement and litigation after the relevant investment or design decision. The costs fall not only on firms, but on research organisations, public bodies, smaller market entrants and individuals who encounter repetitive consent rituals rather than intelligible protection.

6.1 Article 4 and Article 29a: The EDPB Has Changed the Civil-Society Baseline

The constitutional starting point remains sound: the definition of personal data determines the reach of Article 8 of the Charter and cannot be rewritten casually. Yet *EDPS v SRB*[136] and Guidelines 02/2026[137] now point in the same direction. Identifiability depends on the relevant actor, context and realistic means. Pseudonymised information may remain personal for the key-holder while becoming anonymous for a constrained recipient. A wide concept of 'any information' does not remove the separate identifiability threshold.

The strongest civil-society objections therefore address circumvention rather than contextuality. A controller should not escape the GDPR by outsourcing identification, interposing an affiliate, appointing a processor without the key or transferring data to a recipient chosen precisely because that recipient can reconnect it. The July Guidelines expressly support those guardrails. They deny a processor a separate perspective and treat realistic onward transfer as part of the assessment.

The appropriate legislative response is now clearer than either the Commission proposal or the Cyprus retreat. Article 4 should state the contextual test positively and place objective factors, indirect access, anti-circumvention and disclosure responsibility in operative law. Article 29a should become an assurance provision that connects an approved methodology to a rebuttable presumption. The EDPB should update evidence and methods, but the legislature should determine the legal consequence.

---

[131] EDPB, Guidelines 02/2026, which distinguish a contextual approach reflecting the full legal standard from a simplified approach that may treat information as personal even where it is anonymous in the relevant context.

[132] See, among others, noyb, Digital Omnibus Report (2026); European Digital Rights, public materials opposing the Commission's proposed Article 4 clarification; BEUC, 'Protecting EU data and privacy rights in the Digital Omnibus' (24 February 2026).

[133] EDPB, Guidelines 02/2026, online-advertising example concerning pseudonyms derived from device fingerprints and used to link behaviour across services, profile users and select advertisements.

[134] EDPB, Guidelines 02/2026, examples distinguishing a processor acting on a controller's behalf from an independent controller such as a research institute. The processor does not obtain a separate anonymity perspective merely because it lacks the key.

[135] EDPB, Guidelines 02/2026, discussion of legal and contractual constraints. A contract does not by itself create anonymity, but reliable, verifiable and enforceable restrictions may affect which means are reasonably likely when combined with effective technical and organisational controls.

[136] *Case C-413/23 P European Data Protection Supervisor v Single Resolution Board EU*:C:2025:645, especially paras 82–100.

[137] European Data Protection Board, Guidelines 02/2026 on Anonymisation, version 1.0, adopted for public consultation on 7 July 2026, especially the Executive Summary and sections 2.2, 3 and 4.

The Guidelines therefore operate as a two-edged authority. Civil society cannot ask the legislature to defer to the Board and then disregard the Board's acceptance of entity-relative anonymity. At the same time, the 'cannot be ruled out' threshold, the speculative actor set and the indeterminate three-criteria framework show why the legislature should not incorporate the Guidelines by reference. The defensible position is bounded contextuality: an operative actor-specific test, objective factors, anti-circumvention rules and a rebuttable assurance route.

The threshold also rejects a zero-risk standard. *Breyer*, *SRB* and the Guidelines ask whether identification is reasonably likely, not whether an imaginative adversary could construct any conceivable attack. Cost, time, labour, legal access, available technology and realistic chains of cooperation matter because Recital 26 regulates practical identifiability rather than metaphysical possibility.

A burden-of-demonstration rule can address information asymmetry without reverting to universal personal-data status. The actor claiming non-identifiability should document the data, recipient, purposes, auxiliary information, access controls, testing, contracts, circulation boundary and reassessment triggers. Once that evidence establishes an insignificant likelihood, the authority should identify the realistic route that defeats the claim rather than rely on abstract possibility.[138]

This structure aligns incentives with data minimisation. An organisation that retains a key, designs cross-service tracking or can procure identification remains inside the GDPR. An independent recipient that accepts technical and legal constraints, cannot attribute records and submits to continuing assurance can receive a narrower status. The law rewards removal of capability rather than declarations of good intent.

The EDPB's concession should therefore be used directly against the absolute civil-society line: the regulator whose authority opponents invoked as the preferable alternative to legislation has rejected the proposition that anonymity must be universal across all actors. The remaining debate concerns safeguards, evidence and institutional allocation. That is precisely the debate the Omnibus should now settle.

6.2 Article 9: Inference, Residual Sensitivity and the Risk of an Internal Super-Trigger

The status quo case on Article 9 is stronger than caricatures suggest. The case law outlined in the previous section[139] has already rejected the idea that controllers may evade special-category protection simply by avoiding explicit collection while deriving the same sensitive conclusions from other inputs. The structural problem lies elsewhere. If every datum that could contribute, somewhere in an inferential chain, to a sensitive conclusion is treated as special-category data, Article 9 ceases to function as an internal boundary and starts to operate as a super-trigger. The EDPB's own Guidelines 8/2020 capture both sides of that tension. They state that assumptions or inferences regarding special-category data may themselves constitute special-category data, yet also acknowledge that large-scale processing does not automatically fall under Article 9 where the processing does not, in fact, infer such categories and measures are taken to prevent that outcome or its use for targeting.[140] That is the key limiting principle the status quo often understates: Article 9 should attach where the processing is substantively revealing, operationalising or acting upon protected characteristics, not merely because some residual sensitive meaning could be extracted in theory.

This matters acutely for AI development. The Omnibus proposes a new Article 9(2)(k) and paragraph 5 to address incidental and residual processing of special-category data in AI training, testing and validation, coupled with duties to avoid, remove, or, where removal would require disproportionate effort, effectively protect such data.[141] The Joint Opinion does not reject the premise that residual and incidental processing may be unavoidable in some AI pipelines. On the contrary, it acknowledges that this may occur, while insisting that

[138] GDPR art 11(1), which already recognises that a controller need not maintain or acquire additional information solely to identify a data subject where the purposes do not or no longer require identification.

[139] *Case C-184/20 OT v Vyriausioji tarnybinės etikos komisija EU*:C:2022:601; *Case C-252/21 Meta Platforms Inc and Others v Bundeskartellamt EU*:C:2023:537, paras 68–72; Case C-21/23 *Lindenapotheke* EU:C:2024:846, paras 78–84; Case C-446/21 Schrems (Communication of data to the general public) EU:C:2024:834, paras 75–83; EDPB and EDPS, Joint Opinion 2/2026, paras 23–25.

[140] EDPB, Guidelines 8/2020 on the targeting of social media users, Version 2.0, adopted 13 April 2021, paras 123–128.

[141] European Commission, Proposal for a Regulation as regards the simplification of the digital legislative framework (Digital Omnibus), COM(2025) 837 final, recital 33 and art 3(3).

the enacting terms must make clear that the derogation is limited to incidental and residual processing, not to deliberate exploitation or to the prompt collection of sensitive data during deployment.[142] That is the more coherent reform direction: maintain strong protection for deliberate or materially operative sensitive processing, but prevent Article 9 from swallowing ordinary AI development whenever sensitive traits appear only incidentally or probabilistically.

6.3 Article 22: Automated Decision-Making After SCHUFA and Alongside the AI Act

Article 22 illustrates why the status quo cannot simply be equated with the protection of rights. *SCHUFA* confirms that Article 22(1) lays down a prohibition in principle on decisions based solely on automated processing that produce legal or similarly significant effects, subject only to the exceptions in Article 22(2).[143] The Commission nevertheless proposes to replace paragraphs 1 and 2 with a formulation under which such decisions "may" be based solely on automated processing where specified conditions are met, and to clarify that contractual necessity does not require automation to be the only possible decision-making route.[144] The EDPB and EDPS are right to resist any redrafting that risks obscuring Article 22's prohibitory structure and to insist that the provision should remain framed as a prohibition with exceptions and an invocable data-subject right.[145]

Still, preserving the prohibition formula does not solve the real structural problem. Article 22 was drafted for a world in which the main concern was the discrete automated decision at the end of a processing chain. The AI Act regulates something different: risk management, data governance, human oversight, logging and deployer obligations across the lifecycle of high-risk systems.[146] The mismatch is now obvious. The GDPR focuses on a moment of solely automated effect; the AI Act governs system design, deployment and monitoring over time. Without legislative clarification, firms and public bodies will continue to treat the two regimes as overlapping procedural hazards rather than as complementary layers of governance, and many will avoid decision-support systems that are in fact capable of being deployed with meaningful human review and AI Act safeguards.

Article 22 should retain its prohibition while clarifying its relationship with lifecycle-based AI governance. Reform should preserve the ban on solely automated significant decisions unless one of the Article 22 exceptions applies, while making clear that genuinely meaningful human intervention, combined with AI Act risk-management and oversight duties, can keep many decision-support systems outside Article 22(1). The current status quo is under-theorised because it treats a 2016 provision as though it already answers the governance questions posed by 2024–2026 AI legislation. It does not.[147]

6.4 Scientific Research: Fragmentation, Secondary Use and the Cost of National Divergence

The orthodox defence of narrow research exemptions is that they prevent controllers from laundering ordinary commercial processing as "research". That concern is real, but it does not describe the main contemporary problem. EDPB materials on scientific research openly acknowledge that Member State choices on Articles 6 and 9 can have a serious effect on harmonisation, that multi-state health research projects may therefore require heterogeneous legal bases, and that this lack of homogeneity cannot be solved by guidelines or codes of conduct alone.[148] The EDPB's broader studies on secondary use and Article 89 safeguards reach the same conclusion in a different register: while common safeguard types exist, national definitions of scientific research, consent

[142] EDPB and EDPS, Joint Opinion 2/2026, paras 46–52.
[143] Case C-634/21 *SCHUFA Holding AG* (Scoring) EU:C:2023:957, paras 48–55.
[144] European Commission, COM(2025) 837 final, recital 38 and the proposed amendment to GDPR art 22.
[145] EDPB and EDPS, Joint Opinion 2/2026, paras 65–68.
[146] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) [2024] OJ L 2024/1689, arts 9, 10, 14 and 26.
[147] Regulation (EU) 2024/1689 (Artificial Intelligence Act) [2024] OJ L 2024/1689; EDPB and EDPS, Joint Opinion 2/2026, paras 65–68.
[148] EDPB, Document on response to the request from the European Commission for clarifications on the consistent application of the GDPR, focusing on health research, adopted 2 February 2021, paras 13–15.

requirements, and detailed safeguard conditions vary extensively in ways that materially affect secondary use, transparency, and lawful reuse.[149]

The Commission's second GDPR report reinforces the point. It describes the GDPR as a risk-based and proportionate regime, but also recognises that researchers and research organisations need clearer, more actionable guidance and more consistent interpretation across the Union.[150] Since 5 March 2025, the stakes have increased further. The European Health Data Space creates a harmonised framework intended to enable the secure reuse of electronic health data for research, innovation, and public policy.[151] A fragmented GDPR research perimeter now does more than inconvenience researchers. It weakens a flagship European data space strategy that presupposes workable rules on compatible further use, transparency, and safeguards.

The Omnibus therefore moves in the right direction. It clarifies that scientific research may constitute a legitimate interest, that further processing for scientific research is compatible with the initial purpose, and that direct collection cases may also benefit from a disproportionate effort exemption under Article 13(5).[152] Crucially, even the Joint Opinion accepts that the current recitals have not sufficiently clarified whether a separate legal basis is needed for further research processing, and it welcomes the addition of Article 13(5), while recommending that the clarification appear in the enacting terms rather than remain at the recital level. That concession is significant. It amounts to an institutional acknowledgement that the current text does not do enough work.

The cost of inaction is therefore not abstract. Where the boundary of scientific research remains uncertain, actors respond conservatively. They narrow projects, avoid secondary uses, default to consent where it is operationally fragile, and hesitate before entering cross-border collaborations. That effect is especially acute for SMEs, start-ups, hospitals, and smaller research teams that cannot absorb prolonged legal contestation. Fragmentation does not simply generate paperwork. It distorts research design, burdens multi-jurisdictional consortia, and weakens incentives to invest in privacy-preserving infrastructures for lawful reuse. Legislative clarification is therefore not a concession to abuse. It is a precondition for a functioning European research environment and for a Union that wants both robust rights protection and credible data-driven innovation.

6.5 ePrivacy Article 5(3), Terminal Equipment and the Limits of Consent

The cookie debate begins from a legal point that the existing analysis often understates. Article 5(3) of the ePrivacy Directive protects the user's terminal equipment and applies to the storage of, or access to, 'information'. It does not depend on whether the information constitutes personal data. The rule therefore creates an independent terminal-integrity trigger, while the GDPR governs any subsequent processing that concerns an identified or identifiable person.[153]

Planet49 confirms both the breadth and the autonomy of that rule. Consent must result from active behaviour, and the personal or non-personal character of the information stored or accessed does not determine whether Article 5(3) applies. That approach protects the device as part of the user's private sphere. Still, it also means that Article 5(3) reaches technical operations that create little or no data-protection risk unless an exemption applies.[154]

EDPB Guidelines 2/2023 show how far the trigger can extend. The Board treats Article 5(3) as capable of covering cookies, tracking pixels, local processing that instructs a device to disclose information, URL and link parameters, IP-only tracking, unique identifiers and several Internet of Things techniques. The legal problem no

[149] EDPB, Study on the secondary use of personal data in the context of scientific research (1 April 2025); EDPB, Study on the appropriate safeguards under art 89(1) GDPR for the processing of personal data for scientific research (January 2022) 67–68.
[150] European Commission, Communication from the Commission to the European Parliament and the Council, Second Report on the application of the General Data Protection Regulation COM(2024) 357 final, 25 July 2024, 1–2.
[151] Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and amending Directive 2011/24/EU and Regulation (EU) 2024/2847 [2025] OJ L 2025/327.
[152] European Commission, COM(2025) 837 final, recitals 29, 32 and 37 and proposed GDPR art 13(5).
[153] Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector, art 5(3).
[154] Case C-673/17 Bundesverband der Verbraucherzentralen und Verbraucherverbände – Verbraucherzentrale Bundesverband eV v *Planet49* GmbH EU:C:2019:801, especially paras 68–71 and 101–102.

longer concerns cookies alone. A very broad, technologically neutral trigger sits above a narrow and uneven set of exemptions.[155]

The Commission correctly diagnosed consent fatigue but proposed the wrong architecture. Moving personal-data device access into the GDPR while leaving non-personal information in ePrivacy would require actors to classify the same operation before deciding which consent rule applies. The Cyprus text avoided that split by keeping the terminal-equipment trigger in Article 5(3). That choice should be retained, but the exemptions and downstream coordination need a more ambitious rewrite.

The 16 July Council discussion reveals more room for substantive reform than the June compromise suggested. Estonia supported first-party operational analytics and security uses; Poland proposed a two-stage route combining immediate low-risk exemptions with a later transition to privacy-enhancing technologies; France, Czechia and Denmark were open to contextual advertising and broader fraud prevention; Italy supported independent audience measurement; and the Commission expressed openness to all four categories. Spain, the Netherlands, Austria and Hungary demanded tight limits, safeguards or impact evidence. A carefully delimited whitelist therefore has a plausible Council coalition if Ireland separates low-risk first-party operations from cross-service tracking.[156]

The first drafting principle is separation of legal functions. Article 5(3) should decide whether the act of storage or access is permitted. The GDPR should then decide whether subsequent collection, combination, profiling or disclosure involves personal data and, if so, which lawful basis, transparency, minimisation and rights obligations apply. Consent under ePrivacy should not automatically provide a lawful basis for every downstream purpose, and a GDPR lawful basis should not bypass the terminal-equipment rule.

The second principle is a purpose-based low-risk whitelist. Consent should not be required for transmission or a service explicitly requested by the user; security and authentication; service and transaction fraud prevention; tightly constrained audience measurement and analytics; and first-party functions needed to deliver, measure and protect advertising without tracking a person across services. Each exemption should define purpose, retention, combination, disclosure and user-control conditions rather than rely on the vague idea that an activity is merely 'necessary'.[157]

Fraud prevention requires broader drafting than protection of the 'security of the interface'. Fraud often targets the service, transaction, advertiser, merchant or payment flow rather than the user interface itself. The exemption should cover proportionate access needed to detect, prevent and investigate service or transaction fraud, subject to purpose limitation, short retention, access controls and a prohibition on repurposing the signals for behavioural advertising.

Audience measurement should permit a third party acting on the provider's behalf where the arrangement prevents independent use, combination and onward disclosure. The end product should remain aggregated, and the provider should limit retention and access. The law should not demand literal anonymity at the instant of collection where measurement requires transient identifiers; it should demand an architecture that prevents those identifiers from becoming a cross-service profile.

Advertising exemptions require equal precision. Contextual advertising should remain outside consent where selection depends only on the content, application or search context currently viewed and does not draw on behavioural history. First-party delivery, basic measurement, frequency capping and brand safety can also operate without consent where identifiers remain local to the service, retention is short, and no cross-service profile or data-broker linkage occurs. Cross-service tracking, lookalike profiling and behavioural targeting should remain on the consent track.

[155] European Data Protection Board, Guidelines 2/2023 on Technical Scope of Article 5(3) of the ePrivacy Directive, version 2.0, adopted 7 October 2024.

[156] Council working-group discussion summary, 16 July 2026 (on file with the author), section IV, recording the emerging coalition for low-risk ePrivacy exemptions and the safeguards sought by more cautious delegations.

[157] EDPB and EDPS, Joint Opinion 2/2026, paras 96–117. The Opinion supports action against consent fatigue, opposes an unstable split by data type, calls for effective automated choices and questions a media-sector exception.

Browser and operating-system signals can reduce banner fatigue only if the law treats them symmetrically and realistically. A broad refusal signal can operate across services because refusal does not require purpose-specific authorisation. An affirmative consent signal must carry enough information about the controller, purpose and scope to meet the GDPR standard. The legislation should therefore support reusable consent receipts and purpose-specific signals rather than pretending that one global 'yes' can authorise every actor and use.

The media sector needs a sustainable route, but a categorical exemption from reusable preferences would preserve banner fatigue and create a channel for embedded third parties to bypass user choices. The better response is to make the low-risk first-party exemptions usable for publishers while retaining consent for cross-service tracking and profiling. Funding concerns should shape the exemption design; they should not create a general tracking privilege.

Institutional design completes the architecture. Technical standards for browser signals, consent receipts, analytics and first-party advertising should emerge through cooperation among data protection, communications, consumer, competition and cybersecurity authorities. Major standards should carry an evidence note and Innovation Mandate analysis.[158] The EDPB should not determine the full ePrivacy perimeter alone, particularly where the choices affect media pluralism, advertising markets, browser competition and fraud prevention. A legally structured duty to cooperate should govern that cross-regulatory work.[159]

The immediate Omnibus should not pretend to complete ePrivacy reform. Belgium argued for integration with the GDPR, while the Commission reportedly reiterated its intention to repeal the Directive in a later exercise because national transposition and different sanctioning systems perpetuate fragmentation. The practical route is two-stage: amend Article 5(3) now to create a precise low-risk whitelist and interoperable preference architecture; then subject the wider ePrivacy framework to a full, evidence-based fitness review that can decide whether a regulation or GDPR consolidation should replace the Directive.[160]

Figure 6.1. A Coherent Terminal-Equipment and Consent Architecture

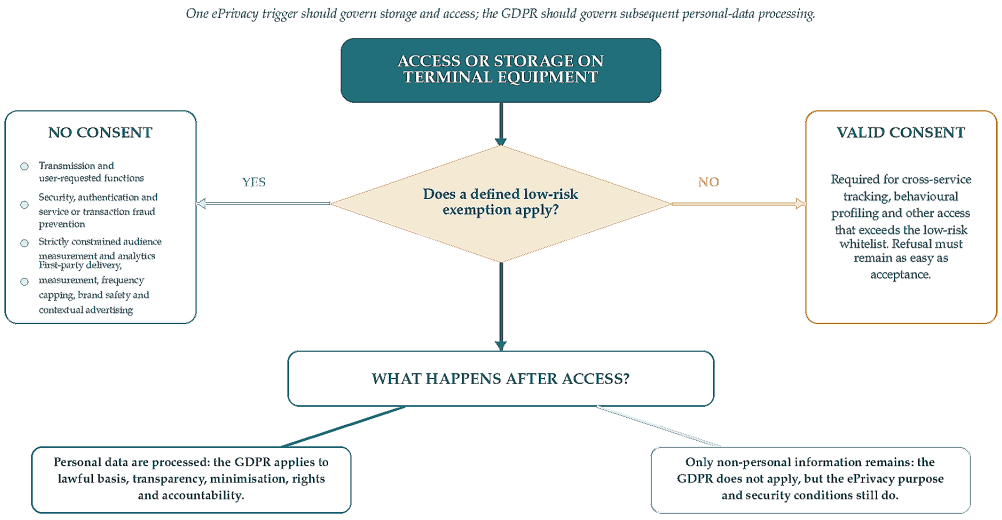

A coherent architecture separates terminal access from downstream data protection and reserves consent for tracking and profiling beyond a defined low-risk whitelist.

[158] M R Leiser, 'The Innovation Mandate: Making Europe's Digital Supervisory State Accountable for Rights, Growth and Scale' (working paper, July 2026) <https://osf.io/preprints/socarxiv/bmcnv_v1> accessed 31 July 2026.

[159] Leiser, 'Recalibrating the EDPB and the EDPS', proposing an enforceable and reviewable duty to cooperate, procedural discipline for major guidance, clearer administrative separation and auditable information governance.

[160] Council working-group discussion summary, 16 July 2026 (on file with the author), Commission response on ePrivacy: the Commission criticised fragmentation caused by national transposition and different sanctions, referred to a future repeal of the Directive and supported continued work on proportionate exemptions for fraud prevention, security, audience measurement and contextual advertising.

Table 6.1. Purpose-Based Terminal-Equipment Rules and Minimum Safeguards

Purpose or technique	Default rule	Minimum safeguards
Transmission and user-requested function	Exempt	Strict necessity for the requested function; no unrelated reuse.
Security, authentication and service or transaction fraud prevention	Exempt	Purpose limitation, short retention, access control and no advertising reuse.
Audience measurement and analytics	Exempt under conditions	First party or processor acting solely on behalf; no independent combination; aggregated reporting.
First-party ad delivery, basic measurement, frequency capping, brand safety and contextual advertising	Exempt under conditions	No cross-service profile, no broker linkage, local or short-lived identifiers and limited retention.
Cross-service tracking, behavioural profiling and personalised advertising	Consent required	Specific, informed and freely given consent; withdrawal and refusal remain easy.
Browser and operating-system signals	Reusable choice mechanism	Broad refusal may be global; affirmative consent must retain granularity for controller, purpose, and scope.

6.6 AI Training, Legitimate Interests and the Deletion of Operative Article 88c

The claim that existing general clauses already provide enough certainty for AI training understates the administrative reality. The Commission proposed Article 88c to confirm that AI development and operation may, where appropriate, pursue a legitimate interest subject to safeguards. Cyprus deleted the operative provision and left a weaker recital. EDPB Opinion 28/2024 accepts legitimate interest in principle but leaves controllers to case-specific balancing. The result is not a settled prohibition or permission; it is a central economic activity governed through non-binding interpretation and national enforcement.

The Council debate also shows that operative clarification remains politically available. Germany, Poland, Denmark, Austria, Czechia and the Commission favoured Article 88c or another binding provision; Belgium, the Netherlands, Malta and Cyprus preferred recital 33a; Finland and Spain accepted a compromise. The viable middle route is narrower than the Commission's original text but stronger than a recital: confirm that AI development may pursue a legitimate interest without creating an automatic legal basis or a hierarchy among Article 6 grounds, legislate the safeguards and retain the ordinary necessity and balancing tests.[161]

That structure distributes uncertainty unevenly. Large actors can retain specialist teams, litigate and maintain alternative datasets. Smaller developers, researchers and downstream providers face a different decision: delay, abandon the use, move the activity or accept a legal risk that investors cannot price. A nominally neutral general clause can therefore entrench incumbents even where it remains formally available to everyone.

A better Article 88c should not create a blanket legal basis. It should confirm that qualifying AI development can pursue a legitimate interest and then specify the safeguards that make reliance credible: purpose definition, source-selection rules, data minimisation, enhanced transparency, respect for machine-readable objections, protection against memorisation and disclosure, special treatment for children and vulnerable persons, and consequences for unlawful source data. Supervisors and courts would retain the three-part Article 6(1)(f) analysis; they could not treat the legitimacy of AI development itself as perpetually unresolved.

[161] Council working-group discussion summary, 16 July 2026 (on file with the author), section II, recording the division between an operative Article 88c and recital 33a and the Commission's preference for binding operative clarification.

The same pattern now runs across the file. The EDPB has accepted contextual anonymity while the law lacks a reliance route. Article 9 needs a narrow response to incidental and residual sensitivity. Article 22 needs a coherent boundary with lifecycle AI governance. Research remains fragmented. Article 5(3) needs a purpose-based exemption architecture rather than repetitive consent. AI training needs an operative pathway rather than recital language. The cost of inaction is continued perimeter-setting through guidance and after-the-fact enforcement.

Table 6.2. Key Pressure Points in the GDPR Simplification Debate

Issue or article	Status quo argument	Structural problem	Cost of inaction	Reform direction
Art. 4 + proposed Art. 29a	Preserve a broad definition and rely on EDPB guidance rather than scope-affecting Commission implementing acts.	SRB and Guidelines 02/2026 now accept contextuality, but the law lacks an operative trigger, guardrails and reliance consequences.	Highest-standard compliance, chilled sharing, divergent DPA practice and no reliable cross-border assurance.	Codify contextuality and anti-circumvention in Article 4; use Article 29a for evidence, assurance and rebuttable reliance.
EDPB institutional role	The Board should ensure consistency through opinions, guidelines and recommendations.	Soft law coordinates but cannot alone create a binding, reliance-worthy and reviewable trigger across the Union.	Practical scope is shaped by expectations that are influential but difficult to contest directly.	Reconstruct the EDPB and EDPS around evidence, cooperation, separation, reasons and review within a statutory perimeter.
Art. 9 – inferential problem	Capture processing that reveals or operationalises protected characteristics and prevents circumvention.	If any inferential link suffices, Article 9 becomes a super-trigger; incidental residual data in AI pipelines remains difficult to manage.	Unpredictable classification, selective enforcement and unstable bias testing or model development.	Protect deliberate sensitive processing strongly; confine Article 9(2)(k) to incidental and residual data with lifecycle safeguards.
Art. 22 – automated decisions	Preserve a prohibition with narrow exceptions and effective data-subject rights.	Article 22 focuses on the decision event; the AI Act governs lifecycle risk, oversight and evidence.	Formal classification disputes, duplicated governance and underuse of meaningful human review.	Keep the prohibition; clarify meaningful intervention and align with AI Act governance without displacement.
Scientific research	Avoid a commercial catch-all and preserve Article 89 safeguards.	National definitions, legal bases and transparency exceptions remain fragmented.	Cross-border studies stall; consortia duplicate legal design; data spaces lose scale.	Define research and compatible further use in operative text with harmonised minimum safeguards.
ePrivacy 5(3) + proposed GDPR Art. 88a	Article 5(3) protects terminal equipment independently of personal-data status.	A broad technical trigger and narrow exemptions make consent the default for low-risk technical functions.	Consent fatigue, poor usability and disproportionate burden for SMEs, publishers and users.	Keep one ePrivacy trigger; add a purpose-based whitelist, reusable signals and GDPR sequencing for downstream processing.
AI training lawful basis	General GDPR bases already apply; AI-specific text must not dilute Article 6(1)(f) or Article 21.	The Council deletes operative Article 88c and leaves recital signals plus case-by-case EDPB interpretation.	Dataset curation and investment remain exposed to *ex ante* uncertainty; larger actors can absorb it more easily.	Create an operative, safeguarded, legitimate-interest pathway with transparency, objection, minimisation and output controls.

Section 7: Systemic Coherence Across the EU Digital *Acquis*

The GDPR no longer operates as a self-contained privacy instrument. It now sits within a broader and increasingly interdependent digital regulatory architecture that includes the AI Act.[162]The Data Act[163], the Data Governance Act[164], the European Health Data Space[165] and the wider programme of common European data spaces.[166] The European Strategy for Data[167] and the Commission's more recent Data Union Strategy[168] both assume that the Union can increase data availability, promote innovation and preserve fundamental rights at the same time.[169] That assumption depends on stable trigger conditions for data protection law. Where the threshold question—whether information is personal, mixed, pseudonymised or effectively anonymised—remains unstable, uncertainty propagates across the wider digital *acquis* rather than remaining confined to GDPR compliance.[170]

In that environment, the definition of "personal data" performs a systemic function. It is not merely a threshold internal to the GDPR. It operates as a switching rule across the wider *acquis*, determining whether other EU instruments apply on their own, apply cumulatively with the GDPR, or operate subject to a full data protection overlay. As Finck and Pallas argued before the current proliferation of EU data legislation, the distinction between personal and non-personal data is already difficult to sustain in complex digital environments.[171] After the Court's judgment in *EDPS v SRB*, that distinction is more clearly contextual and actor-relative: the same dataset may be personal data for one actor, yet not for another, depending on the means reasonably likely to be used for identification in the specific processing context.

That actor-relative approach is doctrinally coherent. It reflects the logic of Recital 26 GDPR and avoids treating identifiability as an abstract or universal property of data. But precisely because it is contextual, it increases the need for legislative clarity across the wider *acquis*. The Commission's 2019 guidance on the interaction between the GDPR and the Free Flow of Non-Personal Data Regulation made it plain that most real-world datasets are mixed, containing both personal and non-personal elements.[172] If that is so, instability at the GDPR threshold is no longer a purely technical privacy issue. It affects the practical operation of data access rights, data-sharing duties, secure processing environments, interoperability requirements, and public-sector re-use regimes throughout the broader EU framework.[173]

7.1 The GDPR Trigger as a Cross-*Acquis* Switching Rule

As a result, the personal-data trigger now operates as a cross-regime switching rule. The AI Act affects the lawful availability of training, validation, and deployment data. It determines whether AI lifecycle obligations must be read alongside GDPR duties on lawful basis, transparency, data-subject rights and special-category data. The

[162] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) [2024] OJ L 2024/1689.

[163] Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act) [2023] OJ L 2023/2854.

[164] Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act) [2022] OJ L 152/1.

[165] Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and amending Directive 2011/24/EU [2025] OJ L 2025/327.

[166] European Commission, 'Appendix to the Communication "A European strategy for data"' COM(2020) 66 final, 'Common European data spaces in strategic sectors and domains of public interest' (19 February 2020).

[167] European Commission, Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, 'A European strategy for data' COM(2020) 66 final (19 February 2020).

[168] European Commission, Communication from the Commission to the European Parliament and the Council, 'Data Union Strategy: Unlocking Data for AI' COM(2025) 835 final (19 November 2025).

[169] European Commission, 'A European strategy for data' COM(2020) 66 final, 1, 4–5; European Commission, 'Data Union Strategy: Unlocking Data for AI' COM(2025) 835 final, 1–2, 18–19.

[170] European Commission, 'A European strategy for data' COM(2020) 66 final; European Commission, 'Data Union Strategy: Unlocking Data for AI' COM(2025) 835 final.

[171] Michèle Finck and Frank Pallas, 'They who must not be identified—distinguishing personal from non-personal data under the GDPR' (2020) 10(1) International Data Privacy Law 11.

[172] European Commission, 'Guidance on the Regulation on a framework for the free flow of non-personal data in the European Union' COM(2019) 250 final, 8–10.

[173] European Commission, 'Guidance on the Regulation on a framework for the free flow of non-personal data in the European Union' COM(2019) 250 final, 8–10.

Data Act determines whether access and portability duties operate subject to the GDPR's full constraints or in relation to non-personal or mixed data. In the European Health Data Space and common data spaces, it affects whether secure processing environments, secondary-use permits and anonymisation requirements are sufficient. In ePrivacy-style terminal-equipment rules, it influences whether choice architecture is governed through the GDPR's consent and objection structure. The same question, therefore, repeats across regimes: personal for whom, in what context, and by reference to which means are reasonably likely to be used?[174]

That is why the Omnibus cannot be understood as a privacy-only file. It is a regime-coherence file. A controller, recipient, researcher, AI provider, data holder or data intermediary does not experience the *acquis* as a set of neat doctrinal silos. It experiences it as cumulative classification work. If the same dataset must be reclassified separately under the GDPR, the AI Act, the Data Act, EHDS rules and sectoral data-space rules, uncertainty is multiplied rather than solved. At scale, the cost is not merely legal advice. It is delayed contracting, a narrower research design, a more conservative architecture, duplicated governance, and a tendency to comply with the strictest plausible interpretation, even when a better-calibrated reading is available.

Figure 7.1. The GDPR Trigger as the Switching Rule of the EU Digital *Acquis*

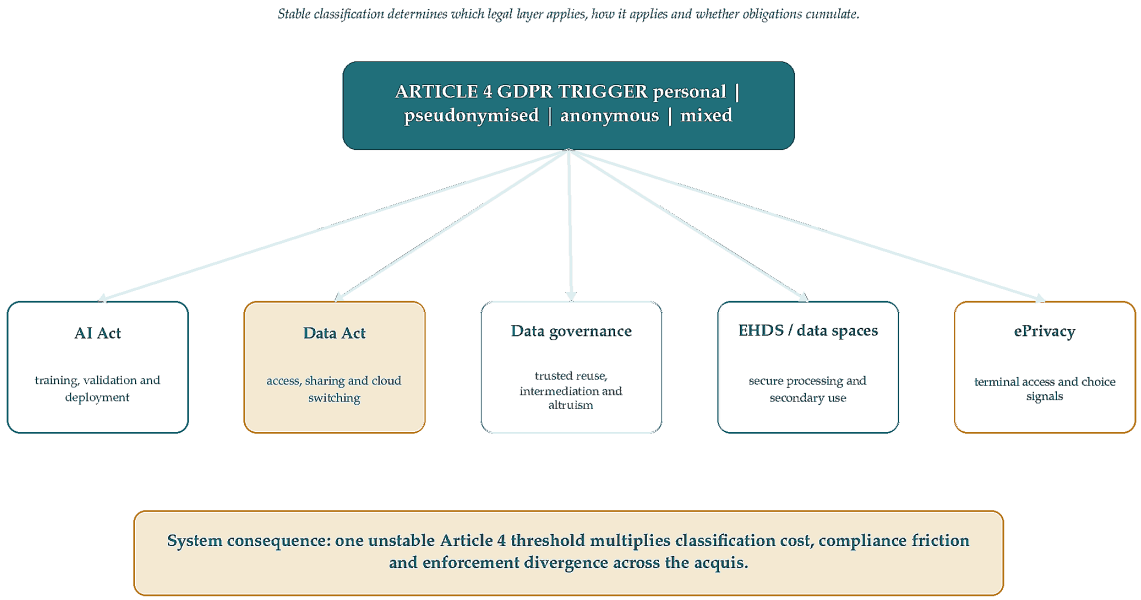

Note: Article 4 GDPR does not merely define a term. In the post-SRB acquis, it decides which legal layer applies, how it applies, and whether adjacent regimes operate cumulatively or independently.

7.2 Mixed Datasets and the Fallacy of Single-Status Data

A significant part of the problem is that the *acquis* still often speaks as though data have a single legal status. In operational reality, many datasets are mixed, relational and mutable. Industrial datasets may contain device data, telemetry, metadata and employee identifiers. Health-research datasets may be pseudonymised for one actor, effectively anonymous for another, and re-identifiable in an uncontrolled public environment. AI training corpora may include ordinary personal data, residual special-category data, public material, synthetic data and non-personal technical information in the same pipeline. The legal classification is not an intrinsic property of the file. It depends on actors, means, controls, purposes and access environments.[175]

Actor-relative identifiability gives the law a way to describe this reality without abandoning protection. It allows the GDPR to attach to the actor that can identify, or can realistically procure identification, while recognising that a constrained recipient may not be in the same position. But that logic must be operationalised. Otherwise, the practical market response is not nuanced classification; it is generalised caution. The data holder treats the

[174] European Commission, Digital Omnibus Regulation Proposal, COM(2025) 837 final, Explanatory Memorandum, 19 November 2025.
[175] European Commission, Guidance on the Regulation on a framework for the free flow of non-personal data in the European Union, COM(2019) 250 final, 29 May 2019, 8–10.

entire dataset as personal data for all actors across all stages because it is the easiest way to avoid enforcement risk. The cost is borne by data spaces, research consortia, start-ups and downstream innovators that depend on the legal significance of safeguards such as pseudonymisation, access controls and secure processing environments.

7.3 AI Act Alignment: Lifecycle Governance Depends on Stable Data Classification

The AI Act reinforces rather than displaces this point. Its architecture is lifecycle-based: risk management, data governance, documentation, logging, human oversight, post-market monitoring and fundamental-rights impact assessment operate across the design, development and deployment of AI systems. Those duties are not substitutes for the GDPR, but they are not indifferent to the GDPR scope either. Where personal data are used, GDPR requirements continue to govern lawful basis, transparency, purpose limitation, rights and special-category safeguards. Where data are effectively anonymous to the relevant actor, the GDPR overlay may fall away, leaving the AI Act duties and sectoral rules to perform their own functions.[176]

The Union has already adjusted the AI Act through the separate Digital Omnibus on AI. Regulation (EU) 2026/1744 was published in the Official Journal on 24 July 2026 and entered into force on the third day following publication. It fixes the application of the principal Annex III high-risk duties to 2 December 2027 and the product-linked high-risk duties to 2 August 2028, and inserts a horizontal Article 4a for tightly bounded processing of special-category data for bias detection and correction. The Union is therefore in a phase of implementation correction, not regulatory stasis.[177].

The correct position is not to subordinate the GDPR to the AI Act. It is to ensure that the two instruments speak coherently. Article 22 should preserve the prohibition on solely automated significant decisions unless an exception applies. The AI Act should govern lifecycle risk. Article 4 should determine, with sufficient stability, whether the data layer of that lifecycle is personal data for the actor in question. Without that division of labour, controllers will over-treat decision-support systems as Article 22 cases, under-use meaningful human review, and duplicate governance work instead of allocating safeguards to the point at which they reduce real risk.

7.4 Data Act, Data Spaces and the European Health Data Space

The same structural problem appears in the Data Act and the wider data-space agenda. The Data Act seeks to facilitate access to and use of data, particularly in connected products and related services, while preserving the GDPR where personal data are involved. That model works only if actors can classify data with enough certainty to know when access duties, trade-secret protections, contract rules, and data protection duties interact. A broad but unstable reading of personal data does not simply create more protection. It creates a veto point at the front of data access and sharing negotiations.[178]

The Digital Omnibus also proposes to consolidate parts of the data *acquis* by folding the Data Governance Act, the Open Data Directive and the Free Flow of Non-Personal Data Regulation into a restructured Data Act framework.[179] That consolidation makes the Article 4 trigger more important, not less. A consolidated data framework needs a stable interface with the GDPR. Without one, formal consolidation may reduce the number of legal instruments while leaving the most important practical uncertainty untouched.

The European Health Data Space shows the point in its most rights-sensitive form. EHDS secondary use depends on secure processing environments, permits and restrictions on re-identification, and a workable distinction between identifiable, pseudonymised and anonymised health data. Those mechanisms are valuable precisely because they attempt to preserve both protection and re-use. If actor-relative identifiability is not legislatively

[176] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 2024/1689.

177 Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulation (EU) 2024/1689 and related instruments as regards the simplification of the implementation of harmonised rules on artificial intelligence, OJ L, 24 July 2026.

[178] Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data (Data Act) [2023] OJ L 2023/2854.

[179] European Parliament, Legislative Train Schedule, 'The Digital Omnibus Regulation Proposal', accessed 19 June 2026; European Commission, COM(2025) 837 final.

stabilised, secure environments will be treated either as insufficient everywhere or as sufficient only after fragmented national interpretation. Neither outcome supports rights-effective European health research.[180]

Figure 7.2. From Trigger Uncertainty to *Acquis* Friction

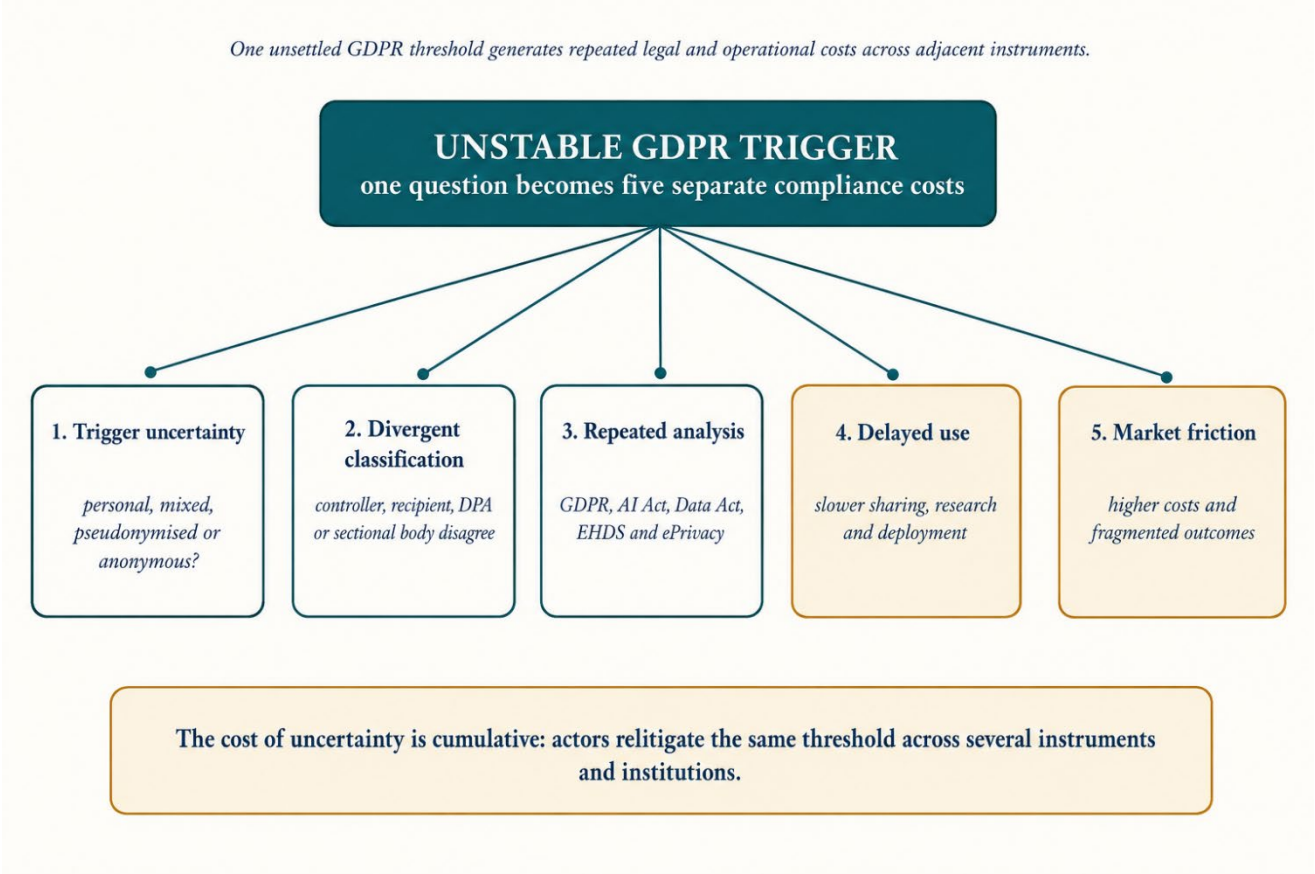

Note: The same threshold uncertainty is paid for repeatedly when a single data architecture must be qualified under several overlapping instruments.

7.5 ePrivacy, Terminal Integrity and the Personal-Data Trigger

ePrivacy creates a different but related switching problem. Article 5(3) applies at the point of storage or access to terminal information, irrespective of whether the information is personal. The GDPR then applies to downstream processing only where the actor handles personal data. A coherent architecture therefore needs two clear questions in sequence: was terminal access exempt or consented, and did the subsequent processing engage the GDPR? Conflating those questions produces either duplicated consent or regulatory gaps.

Contextual identifiability remains relevant after access. Audience measurement, fraud signals or frequency-capping data may be personal for one actor and anonymous for another, depending on the architecture and realistic linkage means. The ePrivacy exemption should define the permitted purpose and device-level safeguards; Article 4 should determine the downstream status. The same actor should not have to relitigate that status across ePrivacy, the GDPR, the Data Act and platform regulation without a portable assurance record.

7.6 What Systemic Coherence Requires

Systemic coherence does not require subordinating the GDPR or ePrivacy to industrial policy. It requires stable trigger logic, clear sequencing and safeguards calibrated to the risk each instrument addresses. The GDPR protects individuals in relation to personal-data processing. Article 5(3) protects terminal integrity. The AI Act governs system risk across the lifecycle. The Data Act and data spaces organise access and reuse. Each regime can remain demanding without duplicating the others' functions.

[180] Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space [2025] OJ L 2025/327.

A coherence test should ask whether the Omnibus makes Article 4 legible across the *acquis*; whether it preserves responsibility for realistic downstream identification; whether terminal-equipment exemptions distinguish low-risk functions from cross-service tracking; whether research and AI rules provide operational pathways with safeguards; and whether supervisory institutions must explain how their guidance interacts with other regulators' mandates. A package that cannot answer those questions will simplify forms while preserving systemic friction.

Section 8: Legislative Settlement — A Targeted Reform Blueprint

The Digital Omnibus should settle trigger logic and supervisory method rather than present itself as a general deregulatory exercise. A clear perimeter can strengthen rights by concentrating obligations where identification and intervention are realistic, rewarding privacy-preserving design, and making non-compliance easier to distinguish from reasonable classification.

The Commission's diagnosis remains sound but incomplete. Article 4, AI, research, cookies and cross-regulatory enforcement now create structural uncertainty. The Cyprus response removed too much operative law and gave too much practical control back to soft law. The July EDPB Guidelines provide useful guardrails but no legal consequence. The Irish reset should therefore support a more ambitious third settlement.

8.1 Article 4: Codify Contextuality and Its Guardrails

Article 4 should state that identifiability is assessed for the relevant entity and processing operation, *ex ante* and *in concreto*, having regard to that entity's actual technical, organisational and legal capabilities and to all means reasonably likely to be used. Objective factors should include cost, time, labour, available technology, access to auxiliary information, legal routes, realistic cooperation and the safeguards that constrain those means.

The enacting terms should also state that singling out does not automatically answer the identifiability question but remains relevant where it enables action directed at a natural person. The assessment should apply separately to record isolation, linkage and meaningful inference. Anonymity need not mean impossibility; the likelihood of identification must be insignificant in reality.

An express anti-circumvention clause should preserve the GDPR where identification can realistically be procured, outsourced or performed by a processor or affiliate; where the controller retains effective control over a key; where separation is artificial; or where a contemplated onward transfer is designed to place the information with an identifying recipient. These rules convert the EDPB's strongest safeguards into binding law.

8.2 Article 29a and the EDPB: Assurance Rather Than Perimeter Substitution

If Article 4 remains politically blocked, Article 29a must expressly govern classification for the whole Regulation rather than operate only as an adjustment to Chapter IV duties. It should not ask the EDPB merely to issue another opinion. It should require a documented assessment, define the evidential categories, provide for recipient-specific and circulation-wide claims, and connect compliance with approved methods, codes or certification to a rebuttable presumption of reasonable classification. The presumption should operate only within the stated actor, purpose, access mode and circulation boundary.

The EDPB should specify testing protocols, attack models, reassessment triggers and sector examples. It should maintain a decision register and update technical methods when the state of the art changes. It should not alter the legal threshold through guidance. A material change in capability, auxiliary data, recipients or attack surface should trigger reassessment and may defeat the presumption.

8.3 A Statutory Assurance Framework for Pseudonymised Data

A usable settlement needs more than a definition. It needs a way for organisations to demonstrate, before sharing or deploying a product, that the statutory test is satisfied. The assessment should identify the relevant actor and processing environment; map direct and indirect access to additional information; test singling out, linkage and inference routes; document legal, technical and organisational barriers; assess onward-transfer scenarios; and record the time horizon and state of the art.

Where that dossier follows an EDPB-approved methodology and is independently certified, the GDPR should provide a rebuttable presumption that the recipient lacked means reasonably likely to identify at the assessment date. This is not immunity. The presumption should not protect concealed facts, artificial separation, collusive arrangements, deliberate re-identification, uncontrolled onward transfer or material changes in capability. It should create reliance for responsible design while preserving enforcement against circumvention.

Articles 42 and 43 already provide the legal framework for certification, seals, and marks. The missing element is a clear connection between assurance and classification. The EDPB can approve common criteria through the consistency mechanism, national bodies can certify against them, and controllers can maintain lifecycle evidence. This approach turns pseudonymisation from a rhetorical safeguard into an auditable compliance and data-sharing infrastructure.[181]

8.4 Article 9: Residual AI Sensitivity with Lifecycle Safeguards

Article 9 should remain a provision of enhanced protection. The reform should not make deliberate sensitive profiling easier. It should solve the narrower problem of incidental and residual special-category data in AI pipelines. The EDPB and EDPS accept that such residual and incidental processing may be unavoidable in some training, testing and validation contexts. The enacting terms should therefore say that Article 9(2)(k) is limited to incidental and residual processing that arises despite appropriate measures to avoid collection or use.[182]

The safeguards should be explicit and lifecycle-based: documented necessity and proportionality; dataset screening; minimisation at source selection; technical controls to prevent sensitive data being used to generate outputs; removal where reasonably practicable; effective protection where removal would require disproportionate effort; logging; internal accountability; and review where the model is repurposed. This is not a loophole. It is a compliance path for a problem that the current Article 9 trigger handles badly.

8.5 Article 22 and the AI Act: Preserve the Prohibition, Clarify the Boundary

Article 22 should remain framed as a prohibition with exceptions. The *SCHUFA* line of authority is important because it prevents controllers from normalising solely automated decisions with legal or similarly significant effects. But Article 22 also needs a clearer relationship with AI Act governance. The reform should distinguish between a solely automated decision and a decision-support system subject to meaningful human intervention. A human rubber stamp is not enough. But a properly designed human review process, with authority to depart from the system output, adequate information, training, documentation and accountability, should keep the system outside Article 22(1).[183]

That clarification would reduce formalistic disputes and focus attention on the substance: whether the person subject to the decision is protected against automation bias, opacity and unchallengeable algorithmic determination. It would also align the GDPR with the AI Act's lifecycle tools without allowing AI Act compliance to displace GDPR rights.

8.6 Scientific Research and Data Spaces

The Omnibus should place the core research clarification in the operative text, not merely in recitals. It should define scientific research broadly enough to include public-interest, academic, health, technological, and innovation-oriented research, while excluding pretextual uses that merely relabel ordinary commercial optimisation. It should confirm compatibility for further processing in genuine scientific research, clarify the relationship between Article 6 legal bases and further processing, and harmonise the minimum safeguards under Article 89.[184]

The point is not to give researchers a free pass. It is to prevent lawful, safeguarded research from being defeated by national fragmentation. The EHDS and common data spaces cannot function if the GDPR research perimeter

[181] GDPR arts 40–43; From Pseudonymised to Anonymous Data (SAMMAN and Team Blaeu, June 2026), section 8.2, recommending codes of conduct, certification frameworks and testing as practical instruments for consistent application.

[182] European Data Protection Board and European Data Protection Supervisor, Joint Opinion 2/2026, paras 46–49; Council of the European Union, Presidency compromise ST 9547/26, 21 May 2026, proposed Article 9(2)(k) and Article 9(5).

[183] Case C-634/21 *SCHUFA Holding AG* (Scoring) EU:C:2023:957; Artificial Intelligence Act, arts 9, 10, 14 and 26.

[184] European Commission, Digital Omnibus Proposal, COM(2025) 837 final, proposed Article 4(38) and amendments to Article 5(1)(b); Council of the European Union, Presidency compromise ST 9547/26, 21 May 2026, recitals 28–32.

remains too uncertain for cross-border use. Harmonisation should therefore focus on minimum safeguards: ethical oversight where appropriate, pseudonymisation or secure environments, transparency adapted to scale, governance logs, access controls, prohibition of re-identification, and enforceable accountability.

8.7 Terminal Equipment and Consent Architecture

Article 5(3) should remain the single immediate rule for storage and access to terminal information, regardless of data status. The legislature should replace the narrow exception structure with a purpose-based whitelist for transmission and user-requested functionality; security, authentication and service or transaction fraud prevention; strictly constrained audience measurement and analytics; and first-party advertising delivery, measurement, frequency capping, brand safety and genuinely contextual advertising. Cross-service tracking and behavioural profiling should continue to require valid consent. The Omnibus should also mandate a subsequent evidence-based review of whether a directly applicable instrument or GDPR consolidation should replace the fragmented Directive.

Browser and operating-system choices should be machine-readable and portable. Services should accept refusal across services without repeated prompts for a defined period. Affirmative consent should require sufficient controller and purpose granularity. The law should support interoperable consent receipts rather than a global opt-in that cannot meet the specificity standard.

The media sector should use the same low-risk first-party routes as other providers, with particular attention to contextual advertising and measurement. A categorical power to ignore reusable preferences should be rejected. Enforcement and standard-setting should involve communications, consumer, competition and cybersecurity authorities under the duty to cooperate.

8.8 AI Training and Legitimate Interests

Deleting operative Article 88c should not leave AI development to recital language and case-specific opinions. A revised and narrower provision should confirm that qualifying development and deployment may pursue a legitimate interest without creating an automatic lawful basis or displacing the hierarchy-neutral structure of Article 6. It should preserve the necessity and balancing tests, enhanced transparency, objection rights, source-selection discipline, minimisation and output safeguards.

Article 88c should also clarify the relationship with Article 9(2)(k). The former supplies a potential Article 6 basis for ordinary personal data; the latter addresses only incidental and residual special-category data subject to avoidance, erasure or effective isolation. Neither provision should authorise deliberate sensitive profiling or remove the controller's duty to assess necessity and proportionality.

8.9 Governance, Review and Institutional Discipline

The Omnibus should impose an enforceable duty to cooperate on the EDPB, EDPS, national data protection authorities, the Commission when enforcing the DSA or DMA, the AI Office, national AI authorities, consumer and competition authorities, cybersecurity bodies and Data Act regulators when their mandates materially overlap. Serious failures to consult, exchange necessary information, record material divergence or coordinate cumulative remedies should become reviewable where they could have influenced the outcome or impaired defence rights.[185]

Major EDPB and EDPS guidance should also trigger an Innovation Mandate within the lawful choice set. A proportionate supervisory impact statement should identify the legal issue, the rights or safety risk, the evidence, affected groups, credible rights-preserving alternatives, distributional effects, cross-border implications and the reasons for the chosen route. A public evidence register, review date and cross-regulatory compatibility note should accompany positions that shape more than one digital regime.[186]

[185] Leiser, 'Recalibrating the EDPB and the EDPS', proposing an enforceable and reviewable duty to cooperate, procedural discipline for major guidance, clearer administrative separation and auditable information governance.

[186] M R Leiser, 'The Innovation Mandate: Making Europe's Digital Supervisory State Accountable for Rights, Growth and Scale' (working paper, July 2026) <https://osf.io/preprints/socarxiv/bmcnv_v1> accessed 31 July 2026.

The Commission should review perimeter effects within three years of application. The review should measure legal certainty, divergence, use of assurance, rights outcomes, SME and research access, cross-regulatory coordination and whether guidance enabled lawful redesign rather than merely added documentation. The review should distinguish legislative defects from administrative implementation choices.

Table 8.1. Legislative Design Principles for a Rights-Preserving Digital Omnibus

Reform area	Legislative settlement	Safeguard against overreach
Article 4	Codify contextual identifiability and the insignificant-likelihood standard.	Include indirect access, processors, affiliates, onward transfer and anti-circumvention.
Article 29a	Create a documented assessment and rebuttable assurance route.	Limit claims by actor, purpose, access mode, circulation and time.
EDPB technical role	Mandate methods, attack testing, examples and a decision register.	Guidance operationalises law and does not move the threshold by stealth.
Certification	Link approved codes and certification to a rebuttable presumption.	No protection for concealed facts, artificial separation or material capability change.
Article 9	Permit only incidental and residual AI-related special-category processing where unavoidable.	Require avoidance, erasure or isolation, documentation and lifecycle controls.
Article 22	Preserve the prohibition on solely automated significant decisions.	Clarify meaningful human intervention and alignment with the AI Act.
Scientific research	Define research and compatible further use in operative text.	Article 89 safeguards, anti-pretext limits and harmonised minimum conditions.
ePrivacy Article 5(3)	Keep one terminal-equipment trigger and adopt a purpose-based low-risk whitelist.	Consent remains for cross-service tracking and profiling; refusal stays easy and reusable.
AI legitimate interests	Create an operative Article 88c pathway for qualifying AI development.	Maintain necessity, balancing, transparency, objection, minimisation and output controls.
Duty to cooperate	Require early notice, lawful exchange, reasons for divergence and remedy coordination.	No consensus duty or inter-regulatory veto; serious failures become reviewable.
Innovation Mandate	Require evidence, alternatives and proportionality reasons for major general supervisory choices.	Applies only within the lawful choice set and never creates a defence to breach.

Conclusion: The Omnibus as Legislative Opening, Not Sufficient Settlement

The political position at 31 July 2026 is more open than it appeared in June, but it has also become more exacting. Parliament remains in an early amendment and impact-assessment phase, and the EDPB has opened its anonymisation consultation. In Council, the 16 July discussion confirmed that Article 4 itself may remain blocked while contextuality, binding criteria, broader low-risk ePrivacy exemptions and an operative AI route retain meaningful support. The 11 September working party will provide the first procedural test of whether Ireland can convert those issue-specific coalitions into a coherent package.

> **Strong data protection does not require the GDPR to attach to every data artefact for every actor whenever identification remains imaginable somewhere. It requires a clear and enforceable perimeter that attaches where identification, intervention and harm are realistic, preserves responsibility for designed disclosure and rewards demonstrable removal of identification capability.**

> **Now that the EDPB has accepted the core contextual principle, the remaining task is to put that principle into law with stronger safeguards and greater legal certainty than either the Commission proposal or the EDPB Guidelines presently provide. Article 29a can serve as the vehicle only if it preserves SRB, rejects speculative downstream contamination, assesses processors by actual capability, supplies objective factors and connects approved assurance to a rebuttable legal consequence.**

The same logic extends beyond anonymisation. Article 9 needs a bounded route for incidental and residual sensitive data in AI pipelines. Article 22 needs a coherent relationship with meaningful human intervention and AI Act governance. Research needs harmonised minimum conditions. Article 5(3) needs an ambitious low-risk whitelist and reusable choice architecture. AI training needs an operative, safeguarded legitimate-interest pathway.

Legislative text alone will not complete the task. The EDPB and EDPS require clearer administrative separation, an enforceable duty to cooperate and procedural discipline proportionate to the market-shaping effects of their guidance. An Innovation Mandate should require evidence, alternatives and public reasons whenever supervisors choose among lawful routes with materially different consequences for rights, certainty and scale. The resulting settlement would not weaken the European model. It would make that model more lawful, more intelligible, and better able to support rights-effective innovation.