A Real Digital Omnibus for European Competitiveness and Security

Robin Berjon · 2026-09-24

supramundane *agency*.

Table of Contents

Introduction

Simplification is a complicated project. The instinct that compelled the Draghi report on competitiveness, the Letta report on the Single Market, and the Niinistö report on civil and military preparedness is sound: that Europe's regulatory architecture has accumulated complexity faster than it has delivered value. Member States echo it with their call for a "clear, simple, smart and innovation- and SME-friendly regulatory framework."[1] The European Commission has consequently proposed a Digital Omnibus "to optimise the application of the digital rulebook."[2]

This report's premise is that good regulation is not an obstacle to European competitiveness but one of its structural preconditions. Businesses, citizens, and public authorities invest, build, and transact within markets to the extent that those markets constitute a *trusted, reliable environment*: where the rules are legible, the protections dependable, the enforcement predictable, and the logic according to which data and value circulate is known in advance. Where that environment exists, all actors can compete. Where it does not, the market defaults to the incumbents whose scale allows them to absorb uncertainty as a cost of doing business. The GDPR, for all its imperfections, has pulled in that direction: it brought us closer to a Single Market for personal data as a trusted environment in which European businesses can build without wondering what rules will apply to them tomorrow.

We must be careful that simplifying our system does not produce a less trustworthy environment. We must structure this evolution around a clear-eyed strategy that applies a detailed understanding of the digital economy to the improvement of European competitiveness. With that objective in mind, this report's primary contribution is in *explaining why* certain changes will produce better outcomes than others.

This document is offered as an expert contribution to ongoing discussions about the data aspects of the Digital Omnibus. Drawn from my experience in digital policy, technology strategy,

[1] European Council meeting (26 June 2025) – Conclusions. https://www.consilium.europa.eu/media/cjtb3oep/20250626-european-council-conclusions-en.pdf

[2] Digital Omnibus Regulation Proposal. https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal

and digital business, I find that the Commission's current Omnibus proposal contains a small number of provisions that achieve genuine simplification and strengthen the reliability of the regulatory environment, alongside a larger number of provisions that, however well-intentioned, would produce outcomes contrary to the legislative aims set out in the Commission's own framing. My purpose is to provide the framework and background information necessary for trilogue interlocutors to identify which provisions can be retained, which can be improved, and which should be reconsidered, so that the final Omnibus achieves the impact that its proponents seek.

Where the proposal succeeds. The introduction of Articles 88a and 88b is the clearest illustration of simplification done well. Article 88a clarifies the rules for access to terminal equipment, narrowing the consent-banner friction that has produced compliance overhead and degraded user experience while delivering limited privacy benefit. Article 88b creates the legal basis for automated privacy signals, an approach that creates a better experience for people, simplified compliance for online services, much better digital economics for ad-driven businesses (especially media), and offers a more effective implementation of fundamental rights than the consent-banner privacy theatre. These provisions move European data protection toward the kind of frictionless, machine-readable interface that the digital economy actually requires.

Where the proposal misfires. The same cannot be said of several other provisions, which on close analysis fail the test their own framing imposes. The proposed amendments to the definition of personal data in Article 4(1), the new legal bases for AI training under Articles 9(2)(k) and 88c, the narrowing of the right of access under Article 12(5), the broad new definition of scientific research, and adjacent changes share a common structural pattern. They do not deliver meaningful compliance simplification for the European businesses that bear the costs of the current regime; they introduce new legal uncertainty that disproportionately burdens small and medium enterprises whilst remaining navigable by entities with sophisticated legal infrastructure; and they reshape data flows in ways that systematically benefit incumbents — predominantly non-European monopolies — at the expense of European challengers. They result in the erosion of exactly the trusted, reliable environment on which European businesses and citizens depend.

Specifically, the Commission's Omnibus proposal in its current form:

- **Fails to deliver meaningful simplification** for the majority of European businesses, whilst simultaneously going well beyond any reasonable definition of "simplification";
- **Benefits monopolistic foreign incumbents over smaller European challengers**, undermining the competitiveness and strategic-autonomy objectives that the Draghi and Letta reports set as the framework's organising aims;
- **Fails to deliver meaningful cost reductions** for businesses: the Commission's own evaluation[3] puts the administrative cost savings from *all* omnibuses (not just the Digital one) at less than €12 billion per year, equivalent to roughly 0.07 per cent of EU GDP — a sum that cannot justify substantive amendments to a fundamental-rights regime;
- **Undermines European security** by facilitating the activities of hostile state and non-state actors. The weakening of the definition of personal data would in particular render large categories of "adtech" data formally non-personal, leaving it fully exploitable by

[3] Simplification. https://commission.europa.eu/law/law-making-process/better-regulation/simplification-and-implementation/simplification_en

malicious entities further down the data-broker chain. Such data has been documented as a vector for tracking sensitive personnel and political leaders, including US soldiers during the Iran war[4] (extensive further examples can be found in the relevant section below), by hostile services. It also contributes to the thriving scam economy, with the elderly, teenagers, and small businesses as its primary victims;

- **Weakens crucial safeguards against systematic online surveillance and behavioural manipulation**, with little corresponding benefit for European businesses beyond a small group of foreign data monopolies.

[4] Mehul Srivastava, Jacob Judah, & James Politi. *US military targeted in Iran war phone-tracking campaign*. 2026. https://www.ft.com/content/44351c74-03c8-45ab-823b-5805c0daca5f

Summary of recommendations

This report does three things: (1) it provides an overarching set of principles to guide discussions on the Digital Omnibus and future reform of EU digital legislation; (2) it identifies aspects of the current Omnibus proposal that would undermine European competitiveness and security; and (3) it proposes targeted improvements. I focus on the data and GDPR-related parts of the Omnibus, because control of and access to data is important to innovation and competitiveness in the digital economy.

In a nutshell, my proposal:

- **Genuinely simplifies compliance** for European businesses, targeting the practices that impose high bureaucratic burden with little data protection benefit.
- **Streamlines citizens' interaction with online services (largely killing banners)**, eliminating most consent prompts while strengthening privacy.
- **Boosts the competitiveness of European businesses and publishers** by protecting their data advantage from the extractive practices of tech monopolies.
- **Bolsters European security** by cutting off the large-scale, data-driven espionage conducted by hostile state and non-state actors, and protecting citizens from the data collection that criminals exploit.
- **Promotes online safety** by reducing the systematic exploitation, surveillance, and manipulation of Europeans, old and young.

My key recommendations:

An efficient legal architecture. Create a Very Large Data Controller (VLDC) category, aligned with DSA/DMA thresholds, concentrating enhanced obligations (strict purpose limitation, separation of data between publicly declared contexts, annual audits, researcher access, and an escalating sanctions ladder) on the largest controllers, while significantly reducing administrative burden for everyone else. Apply core GDPR obligations directly to processors

(Article 28(3)), replacing millions of unenforceable contracts with statutory duties on the firms that hold actual market power. Assign enforcement against VLDCs to the authority of the data subject rather than the country of establishment, ending the de facto subsidy enjoyed by firms established in countries that do not enforce the GDPR. Mandate the EDPB to organise standards for common processing operations with a reformed, OpenStand-compliant European standardisation system, which is arguably the largest untapped source of simplification in the digital acquis. Minimise joint controllership so that the sole-controller-with-processors model becomes the default architecture of the digital economy.

Automation and the data advantage. Retain the current definition of personal data and withdraw the AI-training exemptions, which would legitimise the incumbents' unlawfully accumulated data troves and permanently foreclose European AI competition, and, instead, enforce the law as written against everyone. Keep Articles 88a and 88b as an inseparable package, extend automated signals across the digital acquis (consent, objection, DMA Article 5(2) data combination, minor protection) and to all connected devices, and remove the proposed media exemption, which would lock European publishers into dependency on the tech monopolies rather than protect them.

Security and protection. Reclassify precise geolocation data and data collected by user agents for their own vendors as special-category data under Article 9, closing the loopholes that feed ADINT espionage against European businesses and institutions. Replace covert location tracking for advertising with user-declared coarse location. Protect children by applying automated signals to minors' data with the most protective defaults. Combat the scam economy through deny-lists on vulnerability-based datasets and audiences, advertiser verification duties for VLDCs, and accessory liability for firms and their employees that knowingly profit from the delivery of criminal advertising.

Principles for updating the data protection framework

1. Develop the European data advantage

A basic rule of business: if you know something your competitors don't, you have an advantage. In the digital economy this is the **data advantage**. A business that understands its customers better than its rivals can build better products, and a business that knows its audience well is more attractive to advertisers.

Talk of a "data economy" obscures how this value actually works. Data is not a commodity like oil. Its value depends entirely on what it *refers* to, and for personal audience data that value is **rivalrous**: it diminishes as the data is shared, because each additional holder can monetise it and the originating business loses exclusivity over what made it valuable. When audience data flows from a regular business's service to a large data monopoly that can combine it with data from thousands of other sources and monetise it across search, social media, and display advertising, the smaller business is undermining itself and helping a major competitor.

Consumer-facing businesses hold a natural data advantage: direct relationships with their users, readers, and customers. A publisher knows its audience better than any intermediary; a retailer knows its customers' behaviour better than any platform. This advantage is particularly strong for European businesses who have direct access to one of the world's most valuable audiences. The monopolistic incumbents that dominate digital intermediation have organised markets to undo this advantage: they condition access to advertising demand, app distribution, and analytics on the surrender of first-party data, which they then use to

compete against the businesses that supplied it. Several Omnibus provisions, notably the narrowed definition of personal data and the AI-processing exemptions, would deepen this transfer rather than correct it. The proposed media exemption, though well-intentioned, would do the same by preventing the sector from acting collectively in its own interest to prevent its advantage from being drained to tech monopolies.

The framework should instead protect the data advantage where it naturally sits — with the business the user chose to deal with:

- **Maintain the broad scope of the GDPR**, including its expansive definition of personal data, while reducing burdens for smaller players;
- **Favour the controller of the service the user actually visits**, treating additional parties as data processors in the majority of cases;
- **Require user agents to act in service of their users.** Browsers, operating systems, and AI assistants should not exploit their position to collect data across every service their users touch. When Google Chrome collects the full online behaviour of its users, it is undermining the businesses those users visit;
- **Concentrate obligations and enforcement on the largest controllers**, as the DMA and DSA already do. This creates *transitive compliance*: obligations placed on large platforms propagate contractually to their many business users, shifting burden away from smaller firms.

2. Strengthen European sovereignty and security

Privacy regulation is a national security and public safety asset. The same commercial data flows that erode the data advantage also arm hostile actors, because data broadcast into the advertising ecosystem is available to anyone who wants to listen.[5]

This is not a hypothetical threat. ADINT[6] (Advertising Intelligence) products have been documented, at times offering their customers location histories and behavioural profiles for up to 500 million devices, and have known deployments with intelligence agencies.[7] The same data, that gets unnecessarily broadcast to untraceable parties through advertising networks, identifies vulnerable people — the elderly, teenagers, small businesses — as targets for the fast-growing scam economy.

The industry has long described this data as "anonymous" or “pseudonymous” because it is keyed to advertising identifiers rather than names, but regulators and researchers have established that these identifiers can readily be connected back to clear identities.[8] The Omnibus's proposed subjective definition of personal data would write that fiction into law,

---

5 Johnny Ryan, *Europe's hidden security crisis How data about European defence personnel and political leaders flows to foreign states and non-state actors*. 2024 https://www.iccl.ie/digital-data/europes-hidden-security-crisis/

6 Paul Vines, Franziska Roesner, and Tadayoshi Kohno, *Exploring ADINT: Using Ad Targeting for Surveillance on a Budget — or — How Alice Can Buy Ads to Track Bob*. 2017. https://adint.cs.washington.edu/ADINT.pdf

7 Wolfie Christl, Astrid Perry, Luis Fernando Garcia, Siena Anstis, and Ron Deibert, *Uncovering Webloc An Analysis of Penlink’s Ad-based Geolocation Surveillance Tech*, 2026. https://citizenlab.ca/research/analysis-of-penlinks-ad-based-geolocation-surveillance-tech/

8 The US Federal Trade Commission has clarified that mobile advertising identifiers "offer no anonymity in the marketplace" because businesses regularly link them to names, addresses, and phone numbers.

removing what limited protection we have against hostile intelligence services and organised crime. Europe must leave innocence behind and step up its handling of such dual-use data.

3. Deliver the consistency the GDPR promised

The GDPR was designed as a harmonised framework, with the One-Stop-Shop meant to ensure consistent supervision of businesses operating across Member States. That consistency has not materialised. The Irish Data Protection Commission has only concluded one inquiry into the EU's biggest private-sector data controller, Google, and it took both eight years and concerted pressure from [9] — while Google's European competitors, supervised by other authorities, have been inspected, investigated, and fined in the very markets where they compete with it. This offers a de facto subsidy to foreign businesses established in Ireland, to help them compete against European companies.

What's more, the largest actors are the greatest sources of risk yet their market power equips them to offload their compliance obligations to thousands of smaller partners, as Google does with publishers.[10]

To reflect and correct this centrality, I propose a new category of controller subject to enhanced supervision — the **Very Large Data Controller** — with enforcement assigned to the authority of the data subject rather than the country of establishment.

4. Embrace automation

GDPR compliance today demands extensive manual labour from both sides: data subjects click through banners while controllers process each expression of the same preference as a separate compliance event. This wastes time, and it undermines the law's impact since rights that are costly to exercise go unexercised.

Automation should carry this load, and the GDPR anticipated as much: Article 21(5) already provides that the right to object may be exercised "by automated means using technical specifications". The Global Privacy Control shows what such a signal achieves when it has legal force. GPC was developed from 2020 by a coalition of browser makers, researchers, consumer organisations, and digitally-competent media companies, The New York Times, The Washington Post, and Meredith Digital among them, alongside the publisher trade body Digital Content Next.[11] It learned the lesson of Do Not Track's failure: a signal must be simple, unambiguous, and backed by law. Where it acquired that backing, it worked. Tens of millions of users now transmit the signal, and businesses honour it through a single implementation rather than by processing an endless stream of individual requests.

Publishers helped design GPC because first-party businesses benefit when they can act collectively to protect their data advantage from tech monopolies. Articles 88a and 88b (which are only meaningful used together) bring this architecture into European law.

---

[9] https://www.forbrukerradet.no/siste-nytt/digital/google-fined-e403-million-following-consumer-council-report/

[10] Rebecca Hill. The Register, 2018. *Publishers tell Google: We're not your consent lackeys*. https://www.theregister.com/offbeat/2018/05/01/publishers-tell-google-were-not-your-consent-lackeys/1084029

[11] https://globalprivacycontrol.org/

5. Efficient legal architecture

Simplification rarely works if it eliminates or weakens principles. It should instead make them (a) simpler to apply and comply with, (b) applicable only where necessary, and (c) internally consistent, minimising duplication and contradiction between rules. The current Omnibus does not stick to this approach, and in doing so creates disruption and legal uncertainty. An efficient legal architecture rests on three principles:

- **More risk, more responsibility.** The GDPR treats very different processing operations alike, regardless of how safe or invasive they are in practice. A business broadcasting sensitive data to hundreds of parties should not carry the same responsibilities as a firm conducting analytics alone on comparatively safe data. Compliance burden should track risk.
- **Differentiated controller capabilities.** Businesses with tens of millions of European customers and sizeable compliance departments should not be regulated identically to small firms.
- **Standards.** Most data processing across the vast majority of businesses falls into similar, highly predictable categories (e.g. every newsletter operates in essentially the same way) yet each controller currently performs separate compliance work. Where processing is regular, common standards should carry most of the regulatory burden.

6. Simplification should strengthen protection

Done well, simplification improves rights rather than trading them away since it renders rights more clearly and easily applied. Automated signals make rights exercisable at population scale, standards make protection the default rather than an artefact of individual compliance effort, risk-based allocation directs enforcement to where harm concentrates.

One test of the Omnibus is whether it delivers for the groups the GDPR has so far failed most: children exposed to profiling-driven recommender systems and other data-driven toxic designs, and the vulnerable people targeted by a scam economy supported by brokered data. These failures reflect not protections that are too strong, but protections too complex to exercise and too weakly enforced. By focusing on enforceability and on targeting the best outcomes for the simplest intervention, we can make the internet safer for the most vulnerable.

Proposals for an effective Digital Omnibus

Proposals to keep

This critique of the Digital Omnibus is intended to be constructive. It recognises the value of the simplification exercise and many of its aims, notably reducing overlapping obligations and making compliance more straightforward and less costly.

The proposed Article 88a GDPR, which replaces the ePrivacy Directive's two-step framework with a single harmonised set of rules for accessing and storing personal data on user devices, simplifies the rules for controllers, eases compliance across the Single Market, and improves the consent experience for consumers. Properly implemented and enforced, it will reduce consent fatigue through clear standards while bringing operational simplification for controllers. *However*, Article 88a is *only* viable if it is tied to Article 88b. Without 88b, it will dangerously undermine the data advantage of European companies in favour of incumbent monopolies (as can be seen from Google's active lobbying to keep only 88a).

The proposed Article 88b, enabling the replacement of consent banners with automated signals that inform controllers of users' privacy choices, is welcome and long overdue. It addresses the root cause of cookie fatigue and improves people's ability to genuinely express their preferences. Sections 9 and 10 below set out how to complete it.

The amendment to Article 35 moving responsibility for the deny- and allow-lists of processing operations requiring a data protection impact assessment to the European Data Protection Board harmonises rules across the Single Market and increases legal certainty for controllers operating in several Member States.

The proposed Article 23A in NIS2, creating an EU-wide single entry point consolidating incident-reporting obligations across NIS2, CER, GDPR, DORA, the Cyber Resilience Act, and eIDAS, is exactly the type of simplification the Omnibus should be striving for.

The remainder of this section sets out where the proposal must change, and what should replace it. The proposals are grouped in three parts, following the principles set out above: an efficient legal architecture, automation and the data advantage, and security and protection.

Part I — An efficient legal architecture

These proposals implement the principles of differentiated capabilities, risk-proportionate responsibility, and consistent enforcement: concentrating obligations where scale and risk concentrate, and simplifying everywhere else.

1. Scale-appropriate supervision: the Very Large Data Controller (VLDC)

The GDPR treats a bookshop's newsletter and a monopoly platform's population-scale profiling as the same kind of activity. This is the framework's deepest inefficiency, and the DMA and DSA have already shown the remedy: concentrate obligations on the largest actors, whose processing poses systemic risks and whose market position lets them propagate compliance to thousands of business partners. This creates *transitive compliance*, which shifts burden away from smaller firms while raising standards across the supply chain.

I propose a new category of controller, the **Very Large Data Controller**. In alignment with the thresholds established by the DSA and DMA, a VLDC is a controller with more than 45 million monthly active end users in the Union, or one whose services are integrated by more than 10,000 yearly active business users established in the Union.

VLDCs would face enhanced obligations:

Strengthened Article 5 principles. VLDCs should be held to a demonstrably higher standard of purpose limitation and data minimisation, with narrowly drawn exceptions and meaningful sanctions for non-compliance.

Context separation. The defining feature of the largest controllers is horizontal integration: data gathered in one service — search, mail, maps, social, video, commerce — is combined across all of them, producing profiles no single-service competitor can match and no user ever meaningfully consented to. Purpose limitation already prohibits much of this but it has simply never been enforced at scale. I propose to make it enforceable by design:

- VLDCs must maintain separation between data gathered in distinct contexts, at the physical, operational, and legal level;
- VLDCs must publicly declare their contexts. Article 30 provides the natural instrument: the records of processing of a VLDC should enumerate its contexts and be published, making the silo structure auditable by authorities, researchers, and competitors;
- combining data across contexts requires separate, explicit consent from the user (as defined by DMA Article 5(2)) and, to strengthen enforcement compared to the largely-unenforced DMA 5(2), explicit, publicly-documented approval from the competent supervisory authority;

- each context is treated as a separate data controller, notably for the purposes of Article 88b, so that an automated signal keeps contexts separate by default: the signal then removes any legal basis for cross-context combination.

This makes the digital acquis coherent rather than adding to it. The GDPR's purpose limitation, the DMA's Article 5(2) consent requirement, and Article 88b's automated signals become one architecture, mutually enforcing, instead of three texts policed separately and none effectively. It also strengthens the DMA itself: consent state expressed through a signal is machine-verifiable, refusals persist across interactions, and the DMA's prohibition on re-asking within a year can become auditable.

Impact assessments and audits. VLDCs should carry out data protection impact assessments before any new processing, and undergo mandatory annual audits by the responsible supervisory authorities, focused on the enhanced obligations above.

Researcher access. Mandatory audits should be supplemented by data access for vetted researchers, mirroring Article 40 DSA, focused on systemic risks arising from VLDC processing rather than on personal data.

A sanctions ladder. Sanctions should escalate predictably: fines of up to 10% of global turnover; then periodic penalty payments (daily or weekly) for continued non-compliance; and, where financial penalties demonstrably fail to produce compliance, structural measures as the final rung — criminal liability for senior executives and suspension of market access. The ladder gives supervisory authorities a credible path rather than a single blunt instrument, and gives VLDCs certainty about the consequences of sustained defiance.

Additional VLDC-related considerations are incorporated into the following sections as well.

2. Simplified obligations for other controllers

Below the VLDC threshold, several GDPR obligations produce administrative work out of proportion to any protective benefit. I propose:

Articles 12 and 13 — clarified transparency. Replace the current layered-notice obligations with standardised, machine-readable short-form notices, developed under the standards process set out in Section 5, covering the information data subjects actually use: identity of the controller, purposes, recipients, and how to exercise rights. Full transparency obligations are retained wherever processing involves minors' data or special categories. The aim is clarity: a standard notice that people recognise and machines can read protects better than fifteen pages of bespoke legal prose that nobody reads.

Article 15 — access preserved, boilerplate removed. The right of access must remain intact for all controllers: it is the foundation on which every other right rests. But the obligation to furnish the ancillary catalogue of information (envisaged retention criteria, descriptions of safeguards, and similar) should be removed for non-VLDCs because in practice it does not help the data subject: controllers answer these items with boilerplate. Data subjects are best served by receiving their data, the purposes, and the recipients. Note that this is the opposite of the Omnibus's proposed Article 12(5), which narrows access itself on the basis of the data subject's presumed motives, and which should be rejected.

Article 19. The requirement to notify every recipient of each rectification or erasure is disproportionate for smaller controllers and should be removed for non-VLDCs.

Article 30. Non-VLDCs should maintain a simple list of processing activities without further detail. (For VLDCs, by contrast, Article 30 records become the public register of contexts described in Section 1.)

A note on framing: the case for these simplifications is not that the GDPR crushes European business under compliance costs — survey evidence from data protection professionals shows most obligations are manageable and several are valued.[12] The real competitive injury to European business is described below in Section 4: competing against firms whose establishment in a single Member State has, in practice, exempted them from the enforcement everyone else faces.

3. Reforming Article 28(3): obligations that follow power

The GDPR assumes the controller controls the processor. For the overwhelming majority of European businesses, the reality is often inverted: their processors are a handful of dominant infrastructure providers whose terms are non-negotiable. Surveyed data protection professionals confirm this — 85% report that processors, not controllers, hold the market power, and 80.4% describe processor contracts as hardly enforceable "paper compliance."[13]

The Article 28(3) contractual architecture therefore produces millions of contracts that allocate responsibility to the party least able to bear it, while shielding the party that actually designs the processing. I propose what those professionals overwhelmingly support (84.7%): applying core GDPR obligations (security, data minimisation, documentation, cooperation with supervisory authorities) directly to processors as statutory duties, enforceable against them without the fiction of controller instruction. This eliminates the contractual overhead wholesale (a genuine simplification, of the kind measured in millions of documents), places responsibility with operational power, and ensures a fairer market for commercial services (as per 86.1% of respondents). This provides transitive compliance in action: statutory duties on a handful of dominant processors raise the effective compliance of every European business built on their infrastructure.

4. Consistent enforcement across the EU

Data about enforcement remain poor. EDPB reports based on surveys of national authorities or IMI data are unreliable. This monitoring gap is an Achilles heel of the GDPR. As everyone in industry has noticed: the lead authority for Google has *never* concluded an inquiry into Google under the GDPR, while Google's European competitors, supervised by other authorities, have been inspected, investigated, and at times fined in the very markets where they compete with it.

This, much more than compliance cost, is the regulation's true burden on European business. A rule that binds everyone except the largest players is a subsidy to those players. European businesses today carry the GDPR while competing against firms whose decision to establish in Ireland have placed them beyond its practical reach. Enforcement reform is therefore essential to competitiveness.

---

[12] noyb, *Evidence-Based GDPR Optimisation: Survey of Data Protection Professionals*, March 2026. https://noyb.eu/sites/default/files/2026-03/noyb-Survey_Evidence%20Based%20GDPR-Optimisation.pdf

[13] Ibid.

I propose:

- **For VLDCs, competence follows the data subject.** Enforcement against a VLDC lies with the supervisory authority of the data subject, not the country of the controller's main establishment. This provides a structural correction: it removes the bottleneck without abolishing the one-stop-shop for the thousands of ordinary controllers for whom it works, and it restores to every national authority the jurisdiction over the largest controllers that the current architecture has concentrated in a single Member State with near-nonexistent enforcement. The reform can be understood not as a sanction against any country but as the redistribution of an enforcement load that no single authority could reasonably carry.
- **A monitoring regime.** Article 70 should be strengthened to establish systematic monitoring of all cases involving VLDCs or high-risk processing, including the performance of the relevant supervisory authorities, with published data.
- **An inertia trigger.** Where regulatory inertia persists beyond twelve months, concerned Member States may call on the Commission to intervene and ensure that the competent authority acts.

5. Standardisation

Article 70 obliges the EDPB to "ensure the consistent application" of the GDPR. Its output, while often of respectable quality, is too limited and too abstract to deliver what the regulation expects.

The overwhelming majority of businesses carry out the same kinds of processing. Every bookshop with a newsletter, every school sharing photographs with parents, every employer running payroll performs essentially identical operations. Unfortunately, each starts compliance from scratch, and the advice they buy is inconsistent, needlessly complex, and not infrequently wrong. Standardising the most common processing operations is likely the single largest untapped source of simplification in the entire digital acquis, far larger than any redrafting of the regulation's text, because it determines how the regulation is actually applied.

I propose that Article 70 be amended to mandate that the EDPB organise the production of standards for data processing, in collaboration with European Standardisation Organisations (ESO). In order to avoid such standards being pre-empted by tech monopolies, as many standards organisations have been, it may be necessary to reform the ESO system to include a novel organisation under its umbrella that adheres to the Open Stand principles of due process, broad consensus, transparency, balance, and openness.[14] Such an investment, were it to prove necessary, would serve European sovereignty well beyond data protection, given how much of the digital order is settled in standards bodies.

Standards developed this way offer strong legal certainty (which is what businesses are actually asking for) and propagate through markets on their own. If simplification is to happen across Europe rather than on paper in Brussels, the executive must deploy an intensive effort to standardise and co-regulate.

[14] OpenStand, *The Modern Standards Paradigm — Five Key Principles*, affirmed by IEEE, IETF, W3C, IAB, and the Internet Society. https://open-stand.org/about-us/principles/

Additionally, when surveyed, European data protection officers voiced overwhelming support for changes that increase, rather than decrease, certainty: allow-lists and deny-lists for processing activities.[15] All of this can be delivered by putting Articles 40 and 41 to work, in conjunction with the standards process above. Deny-list entries relevant to security and consumer protection are proposed in Part III.

6. Minimising joint controllership

Joint controllership has become a vector for extracting the data advantage of smaller businesses. In the wake of *Wirtschaftsakademie* and *FashionID*,[16] any embedded tool such as like buttons, tracking scripts, or sign-in widgets can render its host a joint controller with the tech monopoly that supplies it. In theory this allocates responsibility but in practice it allocates liability to the party with no bargaining power and no visibility into the processing, while the supplier writes the joint-controllership terms unilaterally and reuses the data for its own ends. The legal complexity has to be managed by the smaller party while the data advantage is being drained to the larger one.

The goal of reform should be to minimise how often joint controllership arises at all. The default architecture of the digital economy should be a sole controller working with processors, which is the configuration that protects the data advantage and gives data subjects one identifiable counterparty. I propose amending Article 26 so that:

- Joint controllership arises only from the genuine joint determination of purposes and means, evidenced by explicit agreement with meaningful mutual transparency along with by clear and conspicuous joint branding. The mere integration of a third-party tool or service does not suffice.
- A supplier of embedded tools or services is either a processor, where it acts solely under instruction, or an independent controller bearing full and sole responsibility for its self-interested processing, but it cannot use its embedder as a liability shield.
- Where joint controllership does arise, responsibility is allocated according to each party's actual power over the purposes and means, with default liability on the party that designed the data flow.

Part II — Automation and the data advantage

These proposals implement the data advantage and automation principles: keeping data with the controller the user chose, and making rights exercisable at machine speed.

7. Maintaining the definition of personal data

The proposed redefinition of personal data is a mistake, and the mistake is best understood through the data advantage.

The value of information is not intrinsic to the bits: it depends entirely on what the information *refers* to. Data valuation varies completely in its rivalry (how its use by one consumer affects

[15] noyb, *Evidence-Based GDPR Optimisation: Survey of Data Protection Professionals*, March 2026. https://noyb.eu/sites/default/files/2026-03/noyb-Survey_Evidence%20Based%20GDPR-Optimisation.pdf

[16] C-210/16 *Wirtschaftsakademie Schleswig-Holstein* (2018); C-40/17 *Fashion ID* (2019).

the benefits of use by others) depending on what it provides information about. To give some examples:

- The rules of a multi-player game constitute *anti-rival* information: its value increases when they are shared more, since this increases the number of people that the game can be played with.
- The number of neurons in a nematode (302) are *non-rival* information: its value doesn't change as it is shared (the marginal cost of sharing to one additional person is zero).
- A person's purchase intent is *rival* information: if I'm the only party to know about it, I can sell advertising based on that information to advertisers at a high price, but the more that information is shared the less I can monetise it.

The value of audience data is rival information: its value depends on exclusivity. The GDPR's broad, objective definition of personal data is what makes that transfer legally consequential: every entity that touches the data inherits obligations, whatever it claims about its own ability to identify anyone. The proposed subjective standard dissolves this. Data would cease to be personal in the hands of any entity that claims it cannot, or does not aim to, identify the subject. Since the entire third-party data industry — brokers, ad-tech intermediaries, "audience enrichment" vendors — is built on identifiers its participants individually disclaim while collectively resolving, the redefinition converts precisely the extraction layer of the data economy into a GDPR-free zone. Investigative work has repeatedly identified specific individuals from "anonymous" identifier-keyed data[17], and the US Federal Trade Commission has concluded that advertising identifiers offer "no anonymity in the marketplace", since businesses routinely link them to names, addresses, and phone numbers.[18]

Publishers' position illustrates what the redefinition would cost European business. To access advertising demand at viable scale, a European publisher must route its inventory through intermediation infrastructure that requires it to expose audience data (including page context, identifiers, behavioural signals) to the intermediary and, through real-time bidding, to hundreds of downstream parties. The intermediaries, particularly dominant ones, then use that audience knowledge to compete against publishers for the same advertising budgets, offering the publishers' own audiences at lower cost on other inventory. The publisher's knowledge of its readers, accumulated through years of investment, is converted into a targeting asset owned by its largest competitor.

Under the current definition, every entity in that chain carries GDPR obligations, and the publisher retains at least the law's protection against onward exploitation (even if largely unenforced). Under the proposed definition, the data would become legally unprotected. This directly undermines businesses' data advantage, particularly media publishers.

The change would not even deliver certainty in exchange. A definition that varies with each holder's claimed capabilities makes it impossible for a business to assess the compliance of its own providers, worsening what businesses already identify as a genuine pain point. The entities equipped to exploit the ambiguity are those with the legal departments to run interpretive risk at scale. The redefinition thus fails on both of the Omnibus's own criteria: it weakens protection and it complicates compliance, while transferring competitive value from

[17] Ingo Dachwitz & Sebastian Meineck, "Targeting the EU", netzpolitik.org, 4 November 2025 (Databroker Files, with Bayerischer Rundfunk, L'Echo, Le Monde, BNR). https://netzpolitik.org/2025/databroker-files-targeting-the-eu/

[18] FEDERAL TRADE COMMISSION v. KOCHAVA INC., https://www.ftc.gov/system/files/ftc_gov/pdf/26AmendedComplaint%28unsealed%29.pdf

European businesses to the monopolistic incumbents of the tracking economy. Article 4(1) should remain as it stands, together with the settled CJEU case law construing it.[19]

8. AI training: enforce the law as written

The Omnibus proposes to allow AI training on personal data under legitimate interest, to permit processing of special-category data in AI development, and to widen "scientific research" to encompass commercial research (Articles 9(2)(k), 88c). These provisions are presented as support for European AI. Their effect is the opposite.

A legitimate-interest basis for AI training neuters purpose limitation: data collected for any purpose, under any pretext, across any number of services, becomes training matter by whoever holds it.

From a competitiveness perspective, the entities holding the largest accumulations of Europeans' personal data are Google, Meta, Microsoft, and their partners such as OpenAI. They assembled these accumulations over two decades, in significant part through practices that European authorities have repeatedly found unlawful but have rarely prevented. Granting a legitimate-interest basis simply makes it easier for them to keep accumulating more data. European AI companies, which do not possess comparable accumulations, lack comparable consumer channels, and operate in an environment in which digital infrastructure like mobile OSs, browsers, advertising networks, etc. are captured (meaning that incumbents can use them to obtain data), would face the same permissive regime without the means to exploit it and no path to ever building one. This would create an incumbent-favouring regime.

The premise is also false. European AI companies are not held back by data protection law: they are held back by digital markets that monopolists have foreclosed through their thorough control of digital infrastructure from cloud to payments, from search to advertising, from browsers and operating systems to social media and app stores, and ranging over commerce, productivity, mapping, standards, identity, and more.

This is a competition failure that Europe's institutions have the tools to address and yet have chosen not to. The pro-competitive position is the simple one: the law as written applies to everyone, including to the training corpora of the largest firms. If purpose limitation and lawful basis were enforced against the incumbents' data accumulation, the playing field would tilt toward European entrants. Articles 9(2)(k) and 88c should be withdrawn, and the definition of scientific research should retain its established meaning.

9. Expand the surface of automation in Article 88b

Automation in the digital acquis is currently asymmetric: deployed by controllers, denied to data subjects, who must manually express preferences, keep no record of what they have "agreed" to, and have no means to revise those grants. Article 88b begins to correct this for consent banners. The same correction applies far more widely, and the digital acquis already contains the hooks. This is an opportunity to create consistency and improve competitiveness.

[19] noyb, *Digital Omnibus: Analysis of Select GDPR and ePrivacy Proposals*, December 2025, analysis of Article 4(1), surveying *Breyer*, *Nowak*, *Pankki*, *IAB Europe*, and *OC*. https://noyb.eu/en/digital-omnibus-report-v3-analysis-select-gdpr-and-eprivacy-proposals-commission

(a) Consent and objection. Users currently interact with consent-management platforms addressing both ePrivacy terminal-access consent and GDPR legal bases. To update a decision they must clear cookies and re-run the banner. The remedy is a configurable signal mechanism that stores all decisions on the user's device (in the browser, a dedicated application, or the operating system), exercising Article 7(3) withdrawal and the Article 21 right to object, which the GDPR already permits "by automated means using technical specifications" (Article 21(5)).

(b) Gatekeeper data combination. Article 5(2) DMA conditions a gatekeeper's combination or cross-use of data on GDPR-standard consent, and forbids repeating a refused request for a year. A user-side signal makes this more readily enforceable. Combined with the VLDC context separation in Part I, this finally gives purpose limitation, the non-enforcement of which is key to the success of multi-service monopolies, an operational mechanism.

(c) Minors. Article 28(2) DSA prohibits profiling-based advertising to recipients the platform knows with reasonable certainty to be minors. Article 28(1) requires proportionate measures for minors' privacy, safety, and security. Using the GPC signal to this effect would expand the applicability and effectiveness of this provision.

There is no reason to limit automated signals to web browsers, as Article 88b currently appears to do. Signals should operate wherever people access the internet: mobile applications, connected TVs, cars, and beyond.

10. Automated signals without a media exemption

A brief taxonomy explains what is at stake. Running an online service typically involves three kinds of actor: the **primary controller**, operating the service the user chose to visit, several **data processors**, working under its instruction and barred from independent reuse, and **additional controllers**, brought in by the primary controller but free to reuse the data they collect for their own purposes. It is often the case the purposes of additional controllers compete with the primary controller's interests, but the structure of market power is such that the primary controllers have no choice. When only the first two groups are involved the data advantage is protected but when additional controllers enter it is undermined.

This doesn’t only affect publishers: it also impacts advertisers. To give an example, consider a mid-sized direct-to-consumer footwear brand advertising on a large social platform. To measure its campaigns, it installs the platform's tracking script and conversion API, transmitting every product view, basket, and sale. Across thousands of advertisers doing the same, the platform assembles category-level intelligence — which users buy shoes, at what price point, at what cadence — that no individual advertiser possesses. It then sells that intelligence back to the brand's competitors, who bid against it for its own customers. This raises the customer-acquisition costs for *all* footwear brands paying to use the data they provided (and who pass it on to their consumers) and the platform captures the difference. Through this system, advertisers are financing the erosion of their own data advantage. The social platform here is an additional controller as described above: admitted to measure a campaign (something which they could do as a data processor) but free to reuse the data it collects for unrelated purposes.

In a competitive market, no rational actor would admit additional controllers. Publishers and advertisers accept them only because the rules of digital markets are set by tech monopolies,

and access to both advertising revenue for publishers and advertising reach for advertisers is conditioned on surrendering their advantage. This is not a technological requirement: it is a consequence of market power and of the legal architecture that enables it.

A legal architecture that preserves primary controllers and their processors while eliminating additional controllers as much as possible is therefore of direct commercial benefit to online services, particularly media publishers, provided that it applies across the board so that it helps them collectively.

That is precisely the architecture the Global Privacy Control was designed to produce, which is why it was created with the active participation of media organisations such as The New York Times, The Washington Post, and the premium-publisher trade body Digital Content Next (alongside browser makers, researchers, and consumer organisations)[20] and is being finalised as a global standard[21]. GPC automates the exercise of Articles 7(3) and 21 so that the primary controller and its processors operate normally while additional controllers are excluded. The detailed semantics of such global opt-out mechanisms are defined in the W3C's Privacy Principles[22], the standard which defines how data protection works on the web.

The GPC legal architecture is particularly important because not only does it make the exercise of rights easier for people and compliance simpler for businesses, it also makes it possible for publishers and advertisers to keep using the advertising system but to do so in a way that preserves their data advantage and protects it from tech monopolies.

Article 88b must therefore be implemented **without** the proposed exemption for media service providers. The exemption is well-intentioned but precisely backwards. Honouring automated signals, combined with the Article 88a exemptions, places media providers on stronger commercial footing: sole control of their audience data, no banner friction, no leakage to the intermediaries that outcompete them for their own advertisers. Exempting them instead locks them into the hostage position that they already occupy. If the co-legislators wish to support the media sector, the better instrument is a strict application of 88b alongside a well-circumscribed Article 88a exemption for the safest categories of advertising-related processing, which supports advertising revenue without surrendering the audience relationship.

Applied across the board, this architecture eliminates the majority of consent popups and restrains monopoly power in two structural ways. First, the tech monopolies share none of their own first-party data while almost everyone else must share with them: making them lose legal access to a large chunk of the rest of the internet's data curtails that asymmetric advantage. Second, an Article 88b signal works against the horizontal integration of data across a monopoly's unrelated services (maps, mail, search, video, browser, OS, advertising) because it removes the legal basis for combining them, in exact alignment with the VLDC context separation and DMA Article 5(2).

---

[20] https://globalprivacycontrol.org/

[21] Sebastian Zimmeck et al., Global Privacy Control (GPC), W3C Working Draft, 2026. https://www.w3.org/TR/gpc/

[22] Robin Berjon & Jeffrey Yasskin, Privacy Principles, W3C Statement, 15 May 2025. https://www.w3.org/TR/privacy-principles/

Part III — Security and protection

These proposals implement the security and protection principles: no amendment should widen Europe's exfiltration surface, and simplification should strengthen protection for those the current framework fails most.

11. Take the ADINT threat seriously

Online advertising routinely undermines European security by providing ADINT (advertising-based intelligence) to foreign state and non-state actors. Location data is harvested from software development kits embedded in mobile applications (not always with a clear understanding from the application maker) and tied to mobile advertising IDs, supplemented by behavioural and demographic data broadcast through the real-time bidding system. Over time, this produces a detailed dossier.

ADINT has been described in the scholarly literature for nearly a decade: researchers demonstrated in 2017 that $1,000 sufficed to track a target's movements at 8-metre resolution and enumerate their sensitive apps.[23] Since then it has been used to track military and intelligence personnel in Germany[24], to spy on top EU officials, whose home addresses and Berlaymont office locations investigative journalists identified from free data-broker samples[25], to offer targeting of "national security decision makers" and people with chronic diseases through mainstream ad platforms[26], and it is currently deployed at population scale, tracking hundreds of millions of phones daily by ICE in the United States[27]. Using small samples of ADINT data[28], the New York Times was able to reconstitute entire lives[29], identify hundreds of apps that participate in such surveillance[30], monitor protests[31], track children[32], and even follow President Trump around[33]. ADINT has also been an important component of

---

23 Paul Vines, Franziska Roesner, and Tadayoshi Kohno, *Exploring ADINT: Using Ad Targeting for Surveillance on a Budget — or — How Alice Can Buy Ads to Track Bob*. 2017. https://adint.cs.washington.edu/ADINT.pdf

24 Dhruv Mehrotra & Dell Cameron, *Anyone Can Buy Data Tracking US Soldiers and Spies to Nuclear Vaults and Brothels in Germany*. 2024. https://www.wired.com/story/phone-data-us-soldiers-spies-nuclear-germany/

25 Ingo Dachwitz & Sebastian Meineck, "Targeting the EU", netzpolitik.org, 4 November 2025 (Databroker Files, with Bayerischer Rundfunk, L'Echo, Le Monde, BNR). https://netzpolitik.org/2025/databroker-files-targeting-the-eu/

26 Dell Cameron and Dhruv Mehrotra, "Google Ad-Tech Users Can Target National Security 'Decision Makers' and People With Chronic Diseases", *Wired*, 20 February 2025, https://www.wired.com/story/google-dv360-banned-audience-segments-national-security/

27 Joseph Cox, *ICE to Buy Tool that Tracks Locations of Hundreds of Millions of Phones Every Day*. 2025. https://www.404media.co/email/0ba0f6a2-9195-4ced-9c40-92bb72367e7a/

28 Stuart A. Thompson & Charlie Warzel, *Twelve Million Phones, One Dataset, Zero Privacy*. 2019. https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html

29 Jennifer Valentino-DeVries, Natasha Singer, Michael H. Keller and Aaron Krolik, *Your Apps Know Where You Were Last Night, and They're Not Keeping It Secret*. 2018. https://www.nytimes.com/interactive/2018/12/10/business/location-data-privacy-apps.html

30 Stuart A. Thompson & Charlie Warzel, *Smartphones Are Spies. Here's Whom They Report To*. 2019. https://www.nytimes.com/interactive/2019/12/20/opinion/location-tracking-smartphone-marketing.html

31 Stuart A. Thompson & Charlie Warzel, *How Your Phone Betrays Democracy*. 2019. https://www.nytimes.com/interactive/2019/12/21/opinion/location-data-democracy-protests.html

32 Stuart A. Thompson & Charlie Warzel, *Where Even the Children Are Being Tracked*. 2019. https://www.nytimes.com/interactive/2019/12/21/opinion/pasadena-smartphone-spying.html

33 Stuart A. Thompson & Charlie Warzel, *How to Track President Trump*. 2019. https://www.nytimes.com/interactive/2019/12/20/opinion/location-data-national-security.html

Israel's surveillance systems in Lebanon, Iran, and Palestine.[34] It has also been used by Iran in its conflict with the U.S.[35] at a level that has alarmed American lawmakers[36, 37], adding to American frustration with how Europe's lack of GDPR enforcement puts American personnel at risk[38]. Commercial surveillance products built on this substrate are sold worldwide and one such tool offers location histories for up to 500 million devices and has been used by Hungarian domestic intelligence since at least 2022 with vendor documents showing tracking of individuals in Germany, Austria, Italy, Romania, and Hungary[39]. One provider claims to have de-anonymised the advertising identifiers of 95% of Italian mobile devices[40].

Foreign actors need no special access to obtain any of this: the adtech industry, with permissions shaped primarily by Google and Apple, enables this collection in the ordinary course of business.

The Omnibus's proposed redefinition of personal data would write this security failure into law. The entire ADINT supply chain rests on the claim that advertising identifiers are "anonymous" which we know to be incorrect. Under the proposed changes, each intermediary in the chain could assert that the data is not personal, while the chain as a whole delivers to any buyer, including hostile intelligence services, the kind of processing that the GDPR should render difficult: the identification, location, and behavioural profiling of specific Europeans, including military personnel, officials, journalists, and political opposition. Weakening our security at this time would, to say the least, be misunderstanding the geopolitical moment.

In addition to maintaining the current protections, we can help remedy our ADINT vulnerability. At a minimum, **precise geolocation data should be added to the special categories of Article 9** and treated accordingly in enforcement. Location is the primary target of ADINT espionage and supports uniquely powerful inferences about people and organisations. The small number of legitimate actors who genuinely need precise location such as navigation, logistics, or emergency services already take additional precautions. The change imposes essentially no burden on businesses not engaged in dangerous or unlawful practices. (I address coarse location for advertising in the next section.)

It is beyond the scope of this report to make specific proposals for alternative systems that may be built, but the broadcasting of sensitive data by today's real-time bidding systems to

---

34 Omer Benkajob, "Revealed: Israeli Cyber Firms Have Developed an 'Insane' New Spyware Tool. No Defense Exists", *Haaretz*, 14 September 2023,

https://www.haaretz.com/israel-news/2023-09-14/ty-article-magazine/.highlight/revealed-israeli-cyber-firms-developed-an-insane-new-spyware-tool-no-defense-exists/0000018a-93cb-de77-a98f-ffdf2fb60000

35 Mehul Srivastava, Jacob Judah, & James Politi. *US military targeted in Iran war phone-tracking campaign*. 2026. https://www.ft.com/content/44351c74-03c8-45ab-823b-5805c0daca5f

36 Justin Sherman. *Iran War Shows Adversaries Can Exploit Big Data, Too*. 2026. https://www.lawfaremedia.org/article/iran-war-shows-adversaries-can-exploit-big-data--too

37 *Department of Defense Letter to Ron Wyden*. https://www.documentcloud.org/documents/28167310-department-of-defense-letter-to-ron-wyden/

38 Justin Sherman. *Europe's Data Broker Problem Threatens U.S. National Security*. 2025. https://www.lawfaremedia.org/article/europe-s-data-broker-problem-threatens-u.s.-national-security

39 Wolfie Christl, Astrid Perry, Luis Fernando Garcia, Siena Anstis & Ron Deibert, Uncovering Webloc: An Analysis of Penlink's Ad-based Geolocation Surveillance Tech, Citizen Lab Report No. 191, April 2026. https://citizenlab.ca/research/analysis-of-penlinks-ad-based-geolocation-surveillance-tech/

40 Martin Untersinger, "How surveillance companies track smartphone users through advertising data", *Le Monde*, 22 January 2026, https://www.lemonde.fr/en/pixels/article/2026/01/22/how-surveillance-companies-track-smartphone-users-through-advertising-data_6749674_13.html

vast networks of intermediaries, billions of times daily, without security measures[41], is not a necessary feature of advertising technology. It is technically possible to redesign the auction system so that it no longer functions as a broadcast surveillance system, which in turn would shift control away from the intermediaries (who designed the current system in their own interest), and towards publishers and advertisers, while being more respectful of people's rights. It is a path worth investigating.

12. User-declared location

Detailed location emerges from multiple sources (GPS coordinates, proximity mapping against Wi-Fi hotspot databases, IP-address lookup) and users are largely unaware of all of them. Yet legitimate advertising use cases require only *coarse* information: an advertiser selling a product available in a given area needs the area (often the country may suffice), not any precise address or position.

I propose that this coarse, approximate location be declared by the user, in a way that can be exposed by the user agent. As with Article 88b, the Commission should request the drafting of a standard enabling users to declare their coarse location. This should be accompanied by a prohibition on precisely locating a person by any other means **for advertising, profiling, or monetisation purposes**. Locating a person remains lawful where the user has explicitly requested a location-dependent function (navigation, weather, finding nearby services), for emergency services, and where required by law.

13. Data collected by the user agent is sensitive data

User agents (browsers, operating systems, and increasingly AI assistants) occupy a strategic position in the digital architecture: they intermediate all of a person's digital interactions, including the most confidential, from banking to corporate secrets to personal intimacy. A browser observes every page visited and an operating system observes every application used. The W3C's Privacy Principles define the user agent's role as that of a trustworthy agent owing fiduciary duties to its user, "in all situations, including in preference to the user agent's implementer."[42]

When the vendor of a user agent collects data through it for the vendor's own purposes (as Google does with Chrome for instance), the data advantage of *every* other business being visited through this user agent are undermined at once. The processing is high-risk because of the sensitivity of the user data that it produces and it is a competition concern because it converts the agent's privileged position into a corporate espionage apparatus: when a person visits a publisher's site, that visit is part of the publisher's data advantage, but if the browser transmits it to its vendor, the vendor can exploit it to outcompete the publisher in the advertising market. Beyond advertising, this empowers browser vendors to analyse the behaviour of users on their competitors' services.

I therefore propose that **personal data collected by a user agent for its own vendor be added to Article 9 as special-category data**. Only a very small number of firms would

[41] Johnny Ryan, *Europe's hidden security crisis How data about European defence personnel and political leaders flows to foreign states and non-state actors*. 2024 https://www.iccl.ie/digital-data/europes-hidden-security-crisis/

[42] Robin Berjon & Jeffrey Yasskin, Privacy Principles, W3C Statement, 15 May 2025. https://www.w3.org/TR/privacy-principles/

need to deploy more stringent measures in response but the benefits accrue to European businesses and citizens across the board.

14. Child protection

The GDPR's protective promise remains unfulfilled for children, who are exposed to profiling-driven recommender systems that determine, minute by minute, what they see. Articles 8 and 9, read together, already supply the protective logic: processing that profiles a minor's behaviour to infer interests, moods, and vulnerabilities engages special-category data and the conditions on children's consent, and should not operate by default. Enforcement has not delivered this. I recommend the development of strong standards for the processing of minors' data under the standards process of Part I; the extension of Article 88b so that automated signals apply to the processing of minors' data with the most protective settings as the default, and monitoring of enhanced enforcement measures to close the gap between the law's design and children's actual experience.

15. Protect the elderly and vulnerable

The scam economy is now among the fastest-growing forms of organised crime, and it runs on the data infrastructure this report describes. Scholarship increasingly identifies scams not as peripheral abuse but as an organising logic of the contemporary social platform economy, whose defining regularity is that the more vulnerable a person is, the more precisely they can be targeted: children on gaming platforms, teenagers by sextortion, the elderly by investment and impersonation fraud, the indebted by predatory credit[43]. The targeting functions of advertising systems are offered to criminal customers.

The platforms cannot be presumed unwilling hosts. Internal documents reported by Reuters show that Meta projected roughly $16 billion — about 10% of 2024 revenue — from ads for scams and banned goods, showing users on the order of 15 billion higher-risk scam advertisements per day, that its own safety staff estimated **Meta platforms were involved in as much as a third of successful scams in the United States**, that advertisers were banned only above 95% automated certainty of fraud, with enforcement weighed explicitly against "revenue impact", and that the company gamed its own transparency Ad Library against regulators' known search terms while resisting universal advertiser verification[44, 45]. Internal assessments concluded that scam revenue would almost certainly exceed “the cost of any regulatory settlement involving scam ads." They priced in potential fines as a cost of doing business.

I propose:

[43] Lana Swartz, "Social Media in the Scam Age", Social Media + Society, 2026. https://journals.sagepub.com/doi/10.1177/20563051261437922

[44] Reuters, "Meta is earning a fortune on a deluge of fraudulent ads, documents show", 6 November 2025. https://www.reuters.com/investigations/meta-is-earning-fortune-deluge-fraudulent-ads-documents-show-2025-11-06/

[45] Reuters, "Meta created playbook to fend off pressure to crack down on scammers, documents show", 31 December 2025. https://www.reuters.com/investigations/meta-created-playbook-fend-off-pressure-crack-down-scammers-documents-show-2025-12-31/

- **Deny-list entries** (Articles 40/41, per Section 5) prohibiting the sale, sharing, or licensing of datasets that identify individuals by vulnerability category (age-related, financial, health-related, or cognitive) and prohibiting the construction of advertising audiences, including lookalike audiences, from vulnerability signals.
- **Vulnerability inferences as special-category data**: inferred susceptibility to fraud, financial distress, cognitive decline, and equivalent classifications should be treated under Article 9.
- **VLDC obligations on advertising integrity**: as part of the enhanced obligations of Part I, VLDCs operating advertising systems should be required to verify advertisers (KYC), detect and remove fraudulent advertising at thresholds set by the supervisory authority rather than by revenue-protective internal policy, and maintain transparency archives whose integrity is subject to the annual audit, rendered available not through their own sluggish portals but on efficient standardised infrastructure.
- **Structural prevention**: the sole-controller architecture of Part II is itself an anti-fraud measure: keeping data with the service the person chose, and eliminating additional controllers, cuts the supply chain of target lists, behavioural detail, and reach identifiers that industrial-scale fraud purchases from the data economy.
- **Mobilise the organised crime toolbox**: companies that knowingly or negligently support the delivery of scam advertising — which is to say are selling services to criminals, for criminal purposes — are acting as accessories to fraud and must be treated as such, with both the company itself and its employees being held responsible for their organised participation in criminal activity. As things stand, tech monopolies have been granted a de facto exemption to criminal prosecution in fraud cases. This exemption has no reason to exist.

Acknowledgements

This report benefitted hugely from input from the Open Markets Institute, Future of Tech Institute, ICCL, noyb, as well as from media organisations that asked not to be named. This report was independently produced by Supramundane Agency in pursuance of its mission to support democracy and free markets, without external funding. Parties that were consulted in the production of this report do not necessarily endorse its conclusions.

Author

Robin Berjon is a political technologist and entrepreneur with three decades of experience spanning Europe, North America, and Asia. He advises organisations of every size on matters of strategy, technology, and governance. He is Principal at the Supramundane Agency, Deputy Director of the IPFS Foundation, Vice-Chair of the Modal Foundation, and Senior Fellow with both the Future of Tech Institute and Public AI Network.

Previously he was VP of Data Governance at The New York Times (as well as, for a time, VP of Marketing Insights), where he worked on privacy and safeguarding media independence, and Vice-Chair of the board of the W3C. He has also chair multiple W3C working groups and edited numerous standards, including HTML5. His work focuses on building durable democratic governance of technology.

He lives in Brussels. He can be found on the internet:

- **Agency**: https://supramundane.agency/
- **Blog**: https://berjon.com/
- **Microblog**: https://mu.social/profile/robin.berjon.com
- **LinkedIn**: https://www.linkedin.com/in/robinberjon/