



THE DIGITAL OMNIBUS AND ePRIVACY LEGAL CERTAINTY, SCOPE, AND ENFORCEABILITY

EXECUTIVE SUMMARY

The Digital Omnibus proposes a significant reconfiguration of the legal framework governing access to terminal equipment and the confidentiality of communications. While the proposal is presented as a simplification exercise, largely in response to so-called 'consent fatigue' and uneven enforcement, **the chosen legal architecture raises serious concerns. It does not deliver legal certainty**. Instead, it creates parallel and partially overlapping regimes whose boundary depends on unstable and undefined concepts, such as information that 'constitutes or leads to' the processing of personal data. **This approach invites divergent interpretation, litigation, and enforcement fragmentation across the Union.**

More fundamentally, the Omnibus reopens core policy choices concerning the confidentiality of communications and access to terminal equipment. These choices were at the heart of the ePrivacy Regulation negotiations, which sought to modernise and harmonise the framework while preserving an equipment-level rule grounded in confidentiality and device integrity. **By shifting core questions into GDPR lawful-basis logic and broad limitation clauses, while leaving modern tracking infrastructures largely untouched, the Omnibus moves away from that rights-based logic.**

For these reasons, **the proposed changes to ePrivacy should not be pursued through a horizontal simplification package**. Any reconsideration of the rules governing access to terminal equipment and related tracking practices has direct implications for fundamental rights, legal certainty, and enforcement. Such changes require a full legislative proposal and a comprehensive assessment of their impact, and should, if considered at all, be addressed in the context of the Digital Fitness Check.

This paper recognises that **some of the objectives invoked by the proposal, such as reducing manipulative consent practices and enabling more durable expressions of refusal, are legitimate**. However, **the substantive elements associated with proposed Articles 88a and 88b should remain within the ePrivacy framework, not the GDPR, in order to preserve regime specificity and coherent protection**. Any future reform would also require significantly tighter safeguards, including narrow and verifiable limits for consent-free access to terminal equipment and strict conditions for privacy signals, to avoid creating new tracking layers or delaying the effective exercise of rights.

Core recommendations

1. Keep a **single device-access rule inside ePrivacy**, applicable to access to terminal equipment regardless of whether the accessed information is personal data. This preserves prevention at the point of access and avoids boundary disputes.
2. **Relocate the substance of Articles 88a and 88b to ePrivacy**, as amendments to Article 5(3) and related provisions, preserving the distinct logic of confidentiality and terminal equipment integrity.
3. **Tighten consent-free allowances with enforceable technical constraints for audience measurement and security**, including bans on persistent identifiers and fingerprinting under the allowance, strict purpose limitation, short retention, and explicit bans on reuse.
4. **Privacy signals should be defined in law and be mandatory and binding from the outset**. Their legal semantics, scope across web and app contexts, lifecycle symmetry including withdrawal, and verifiability requirements should be fixed in legislation rather than deferred to future standardisation.
5. **Clarify the interface**: ePrivacy governs access to terminal equipment and communications confidentiality. GDPR governs subsequent processing of personal data once collected. This restores a clear sequence and supports predictable enforcement.

1. WHAT THE PROPOSAL CHANGES IN EPRIVACY TERMS

The Omnibus makes three changes that matter for the ePrivacy architecture.

First, it **narrows Article 5(3) ePrivacy with a new sentence**: 'This paragraph shall not apply if the subscriber or user is a natural person, and the information stored or accessed constitutes or leads to the processing of personal data.'

Second, it **introduces a new GDPR device-access regime** for natural persons through Article 88a, plus a GDPR signal regime through Article 88b.

Each of these changes affects scope, competence, and enforceability. Taken together, they **shift the system away from a preventive, equipment-level rule towards a mixed model where access, processing, and limitations are assessed through overlapping and partially inconsistent tests**.

2. THE CORE ISSUE: SCOPE SPLITTING THAT UNDERMINES LEGAL CERTAINTY

The existing legal logic is simple: Article 5(3) ePrivacy governs the act of storing or accessing information on terminal equipment, regardless of whether the information is personal data. The GDPR governs the subsequent processing of personal data, once collected. The Omnibus replaces this functional separation with a data-category **split**. The boundary now turns on whether the stored or accessed information '**constitutes or leads to**' personal data processing.

That boundary is legally and technically unstable.

- '**Leads to**' is **undefined**. It can cover correlation, inference, linking, or later enrichment. Different actors will interpret it differently.

- In real systems, controllers often cannot know at the moment of access whether information will become personal data only after correlation. That makes it **difficult to know which rule applies at the moment of access**.
- The drafting creates a plausible compliance and enforcement gap: information that 'leads to' personal data processing can fall outside Article 5(3) ePrivacy, while Article 88a GDPR applies to the 'processing' of personal data.

This is the opposite of simplification. It produces uncertainty for controllers, uncertainty for people, and uncertainty for authorities, including on which regulator is competent and which complaint route applies.

It also creates **perverse incentives**. If the GDPR regime is more permissive for device access than Article 5(3) ePrivacy, controllers gain incentives to frame the accessed information as personal data, or to quickly convert it into personal data, to benefit from the more permissive track. That runs against minimisation and against coherent upstream limits on device access. The Digital Omnibus proposed redefinition of personal data complicates this situation.

Additionally, effective protection at the point of access requires that **the principle of 'no tracking before consent' be applied in practice across all tracking techniques**, including non-cookie identifiers such as fingerprinting, SDK-based tracking, server-side collection, link decoration, and probabilistic identifiers. Enforcement that focuses primarily on the wording or design of consent banners, rather than on the underlying tracking infrastructure, fails to address the source of the interference and does not provide legal certainty."

3. WHY MOVING DEVICE ACCESS INTO THE GDPR CHANGES THE RIGHTS LOGIC

ePrivacy protects confidentiality of communications and integrity of terminal equipment as an upstream safeguard. **This protection is conceptually linked to secrecy and intrusion prevention, not only to personal data processing.** Placing the main device-access regime in the GDPR shifts the baseline in two ways.

First, it **turns an equipment-level prohibition into a lawful-basis and limitation exercise**, where interference becomes easier to justify through general legal bases and broad limitation grounds.

Second, it **changes how future legislative debates will play out**. Article 88a(2) explicitly opens consent-free access based on Union or Member State law aligned with Article 6 and objectives in Article 23(1) GDPR. Article 23(1) contains a wide list of objectives. This is a materially wider and more flexible opening than the traditional ePrivacy logic, which sits inside a strict confidentiality framework.

This is not an abstract point. Device access rules shape the technical feasibility of tracking, surveillance, and device-level extraction. **Once the system accepts broader openings for access, the barrier moves from prevention to ex post compliance arguments.**

4. ARTICLES 88A AND 88B: OBJECTIVES THAT CAN BE ACCEPTED, BUT NOT IN THIS LEGAL HOME

The proposal contains elements that respond to real problems, especially around manipulative banners and the lack of durable refusal tools.

- A requirement to refuse consent through a 'single-click button or equivalent means' can reduce interface abuse **only if refusal has durable effect and blocks non-essential access to terminal equipment as well as any subsequent processing, reuse, or enrichment of data collected through that access**. Otherwise, refusal merely alters the interface without constraining the underlying data flows.

- Machine-readable signals can reduce banner friction and support enforceable refusal, but only if the law defines their semantics, scope, and governance, and only if implementation does not create new tracking layers.

The problem is the system design: these elements are inserted into the GDPR while the ePrivacy rule is narrowed. That creates two co-existing regimes and new boundary disputes. **A rights-aligned approach would keep a single device-access rule within ePrivacy, modernise it, and then let GDPR regulate subsequent processing.** That is closer to the logic that was pursued during the ePrivacy Regulation negotiations: harmonise and modernise ePrivacy, rather than relocating its core into the GDPR.

Additionally, it's critical to mention that **meaningful consent framework presupposes an upstream structural decision by the controller**. Before any consent request is presented, the controller must already have determined which data flows are strictly necessary to provide the service requested and which are optional. Only the latter may be subject to consent. **Where this separation is not made in advance, consent interfaces cannot ensure meaningful choice, regardless of how they are designed.**

5. TIGHTENING CONSENT-FREE DEVICE ACCESS: WHY PURPOSE LABELS ARE NOT ENOUGH

Article 88a(3) introduces a list of consent-free purposes that the proposal presents as low risk, including first-party audience measurement using aggregated data and security of services or devices. **The concern does not lie in the existence of narrowly framed allowances as such, but in defining them primarily through purpose labels without enforceable technical constraints.**

Where exemptions are framed by reference to purpose, compliance assessment shifts from observable technical acts to internal intent and downstream use. This makes enforcement harder and reduces legal certainty, as the same tracking techniques and infrastructures can support multiple purposes depending on how data are later combined or reused.

Without **clear technical limits at the point of collection**, purpose-based allowances risk enabling broad data collection under compliant narratives, while leaving supervisory authorities to assess legality only *ex post*, with limited visibility into internal processing chains.

Audience measurement

Aggregation must operate as a constraint at collection, not only as a reporting output. Otherwise, raw device-level data can be collected widely, retained, and then later aggregated, while still enabling profiling, correlation, or repurposing.

A workable allowance needs **clear conditions**, such as:

- first-party only, single service context, no cross-site or cross-app measurement
- no persistent identifiers, no fingerprinting, no device graphing
- aggregation at source and short retention of any raw event data
- prohibition of downstream reuse for optimisation, advertising, profiling, or AI training

Security

Security allowances regularly expand into tracking bypasses in practice, because fraud and abuse prevention are used to justify persistent identifiers, fingerprinting, long-lived logs, and server-side matching.

A workable security allowance needs **clear limits**, such as:

- strict linkage to concrete security of the requested service or the device used for that service
- rotating identifiers and short retention
- explicit prohibition of reusing security-derived data for analytics, optimisation, or advertising

Without these conditions, consent-free access becomes an infrastructure for behavioural tracking under different terminology.

The six-month rule

The temporal dimension of consent is also relevant. **Limiting how often a person may be asked again, as proposed in the Omnibus, does not address the persistence of tracking once consent has been given.** A framework in which consent can remain valid indefinitely risks making tracking permanent by default. Limiting the validity of consent over time would better reflect its contextual nature and would contribute more directly to reducing long-term tracking pressure.

6. PRIVACY SIGNALS: STANDARDISATION, TRUST, AND ENFORCEABILITY

PROBLEMS THAT NEED LEGAL ANSWERS

Privacy signals raise four interlinked issues that a law must resolve if signals are to improve legal certainty rather than reduce it. **Automated privacy signals can reduce repeated prompts and support durable refusal only if they are recognised as legally valid expressions of objection or withdrawal.** This includes browser-level signals, such as Global Privacy Control-type mechanisms, where they clearly convey a refusal of non-essential tracking. Without explicit legal recognition of such signals, their effectiveness remains contingent on voluntary compliance and does not reduce legal uncertainty.

Privacy signals must operate across the full technical stack. This requires obligations not only for browsers, but also for operating systems and app environments, since a growing share of tracking takes place in-app through SDKs and through OS-level identifiers and permissions. A signal framework that is browser-only will not provide legal certainty or meaningful protection in practice.

Technical interoperability

Different browsers, apps, and device environments can implement signals differently. **Without clear standards and legal semantics, this produces inconsistent outcomes and forum shopping.**

To reduce fragmentation, the legislation should require operating systems to provide a user-facing, system-level mechanism to express refusal of non-essential tracking, and require that browsers and apps can read and transmit that signal in a standardised way. This should **apply equally to web and app contexts.**

The signal framework should prevent browser or OS vendor lock-in. Operating systems and browsers should be required to allow third-party privacy and consent management tools to generate, manage, and transmit signals on the user's behalf. Otherwise, the actors with the strongest commercial incentives in the tracking ecosystem may become the practical gatekeepers of how refusal is expressed and interpreted.

Trust and attribution without new tracking

Signals must reflect a person's choice without requiring persistent identifiers or device-level markers that create a new tracking layer. If a controller needs stable identifiers to recognise a refusal signal across contexts, the signal becomes an additional tracking mechanism.

Scope and hierarchy

The law needs to clarify where a signal applies. Does a browser signal override service-level prompts. Does a device setting override app-level consent. Is refusal durable, and for how long. Ambiguity invites circumvention and inconsistent compliance.

Enforcement in a server-side world

Authorities need practical ways to test compliance. That is difficult when tracking shifts server-side, through gateways, tag managers, or event ingestion infrastructures. A signal regime without verifiability requirements risks becoming performative.

The Omnibus currently adds further problems:

- **The proposal delays the legal effect of signals.** Controllers are required to respect refusal signals only after 24 months, while browsers and device manufacturers are given up to 48 months to implement them. During this transitional period, Article 88a would already apply, yet individuals would lack an effective technical means to express refusal in a durable and interoperable manner. This temporal mismatch weakens protection precisely at the moment when consent-based access rules are being relaxed, and creates uncertainty for both controllers and supervisory authorities about applicable obligations.
- The proposal **leaves the core legal semantics of signals undefined**, explicitly relying on future technical or regulatory development. This shifts decisive questions away from the legislature and into standardisation bodies or informal coordination processes. Experience with previous initiatives, such as Do Not Track, shows that where legal meaning is not fixed in law, standardisation processes become arenas for delay, dilution, or fragmentation. Signals risk acquiring different meanings depending on implementation choices, market power, or sectoral practices, rather than functioning as a uniform expression of rights across the Union.
- This approach also **raises institutional concerns**. Standardisation bodies are not designed to resolve questions of fundamental rights, proportionality, or scope of legal obligations. Delegating the definition of what constitutes a legally valid refusal, objection, or consent signal to such processes risks displacing normative choices from democratic law-making into technical fora, where accountability and judicial review are limited. From an enforcement perspective, this makes it harder for authorities to assess compliance, as disputes shift from breaches of law to disagreements over standards, versions, and interpretations.
- **The design of the signal regime fragments the lifecycle of rights.** The proposal automates refusal and objection but omits withdrawal pathways, despite the principle that withdrawal must be as easy as giving consent. This creates ambiguity as to what signals legally express and how long their effects last. In practice, controllers, users, and authorities are left to infer whether a signal covers refusal only, objection only, or also withdrawal. Such ambiguity undermines predictability and invites inconsistent application across jurisdictions.
- **The exemption for media service providers** from the obligation to respect refusal signals compounds these problems. Signals are presented as a general mechanism to reduce user burden and enable durable choice. Allowing an entire sector to disregard them undermines that rationale and breaks the universality of the mechanism. It also sets a precedent for sector-specific opt-outs from fundamental rights safeguards, which can easily expand beyond media. From an enforcement perspective, this carve-out complicates compliance assessments, as the legal effect of the same signal would depend on the classification of the actor rather than on the nature of the interference with terminal equipment.
- **The treatment of SMEs raises similar concerns.** While proportionality and administrative burden are legitimate considerations, exempting micro, small, or medium-sized enterprises from technical obligations related to signals risks creating a two-tier system. In such a system, the same technical signal may have different legal effects depending on the size of the controller. This undermines legal certainty for individuals, complicates enforcement for authorities, and incentivises regulatory arbitrage through organisational structuring. Proportionality concerns are more coherently addressed through guidance, phased implementation, or enforcement discretion, rather than by weakening the universality of rights or the binding nature of refusal signals.

If signals are to be part of the future framework, and they should, the law should define at least: semantics, scope by context, lifecycle symmetry including withdrawal, and verifiability duties. Otherwise signals will not reduce legal uncertainty. They will move it into standardisation fights and compliance ambiguity.

7. 'COOKIE FATIGUE' AS A SYSTEMIC PROBLEM, NOT A BANNER PROBLEM

The Digital Omnibus repeatedly invokes '**cookie fatigue**' as a **justification** for loosening device-access rules and expanding consent-free pathways. This **framing is incomplete** and risks misdiagnosing the problem.

So-called cookie fatigue does not result from the existence of Article 5(3) ePrivacy as such. It is the product of three interacting dynamics: pervasive tracking as a default business model, widespread use of deceptive or manipulative interface design, and weak enforcement against unlawful tracking practices. Addressing only the surface manifestation, namely consent banners, without **tackling these underlying drivers will not reduce fatigue in a durable way**.

First, **privacy signals can play a role, but only as part of a broader approach**. Signals can reduce repeated prompts and enable durable refusal, yet they cannot compensate for a regulatory environment in which tracking remains economically attractive and legally easy to justify. If consent refusal does not meaningfully restrict tracking practices, signals become a technical convenience rather than a rights-enabling tool. Their effectiveness depends on clear legal semantics, enforceability, and a regulatory backdrop that actually constrains tracking behaviour.

Second, **deceptive and manipulative design practices are a central contributor to consent fatigue**. Many consent interfaces are designed to steer users towards acceptance, obscure refusal options, or fragment choice across multiple layers. This is not primarily a data protection issue, but a consumer protection one. Addressing such practices through a future Digital Fairness Act, including by treating manipulative consent flows as unfair commercial practices, would directly reduce the pressure placed on users at the moment of choice. Without tackling dark patterns and interface manipulation, changes to consent mechanics alone risk entrenching fatigue rather than alleviating it.

Third, **cookie fatigue cannot be meaningfully addressed without regulating the adtech infrastructure** that generates constant consent requests in the first place. The proliferation of tracking requests is driven by complex, multi-actor advertising ecosystems that rely on continuous data extraction, profiling, and sharing. As long as these infrastructures remain largely intact, users will continue to face repeated choices, regardless of whether tracking uses cookies, SDKs, server-side collection, or other techniques.

This also highlights a broader point: cookies are increasingly no longer the central technical mechanism. **Tracking has shifted towards fingerprinting, probabilistic identifiers, server-side tracking, link decoration, SDK-based data flows, and data brokerage**. Framing the problem around 'cookies' risks anchoring regulation to a declining technology, while leaving newer and less visible forms of tracking insufficiently constrained. A rights-aligned response must therefore address pervasive tracking as a whole, not specific storage techniques.

Finally, **enforcement plays a decisive role**. Weak or inconsistent enforcement against unlawful tracking practices allows non-compliant models to persist and multiply, increasing the number of consent requests users encounter. Effective enforcement, particularly in the adtech space, would reduce the prevalence of unlawful tracking and, as a consequence, reduce the frequency of consent interactions. In this sense, enforcement is not opposed to usability; it is a precondition for it.

Where violations are structural, enforcement responses must also be structural. This includes measures that address the tracking stack itself, such as orders to suspend unlawful tracking infrastructures and sanctions that reflect systemic infringements, rather than reliance on warnings or improvement plans that leave underlying data flows unchanged.

Addressing cookie fatigue therefore requires a multidimensional approach: enforceable privacy signals, robust action against deceptive design through consumer law, structural constraints on adtech infrastructures, and stronger enforcement of existing rules. **Without this broader perspective, reforms risk treating symptoms while leaving the underlying causes untouched.**