Brussels, 15th March 2026

Consultation Statement by the European AI Forum (EAIF) on:

Digital Omnibus on Data

We welcome the opportunity to provide feedback on the European Commission's proposed amendments to the General Data Protection Regulation and related data governance frameworks as part of the Digital Omnibus simplification package.

Europe's ability to develop and deploy artificial intelligence at scale is inseparable from the legal environment in which European AI companies operate. While the United States and China have moved decisively to create regulatory conditions that support large-scale AI training, European startups and scale-ups continue to navigate a framework characterised by legal uncertainty, fragmented national enforcement, and compliance burdens calibrated for large enterprises. The consequence is not merely administrative friction: it is a structural competitive disadvantage that causes promising European AI companies to train models abroad, relocate to more permissive jurisdictions, or delay product development while competitors in other markets advance. European AI sovereignty, understood as the capacity of European institutions, companies, and citizens to rely on AI systems built and governed in Europe, cannot be achieved through regulation alone. It requires companies that can scale within the EU, on a legal foundation that is clear, proportionate, and internationally competitive.

The Digital Omnibus data package is an opportunity to establish that foundation. For European AI startups and SMEs, the current framework creates structural compliance challenges that are disproportionate to the actual risks involved: overbroad interpretations of personal data, unclear legal bases for AI development, fragmented national requirements for impact assessments and breach notification, outdated rules governing consent, and conflicting obligations across the GDPR, the ePrivacy Directive, and the AI Act. These challenges generate significant legal uncertainty and compliance costs, and actively constrain the responsible use of data for AI development, research, and innovation.

This paper sets out the European AI Forum's (EAIF) detailed assessment of the Commission's proposed GDPR and data governance amendments. For each provision, we evaluate whether the proposed changes adequately address existing legal uncertainty and proportionality concerns, identify where careful implementation will be critical, and highlight where the balance between innovation facilitation and fundamental rights protection requires particular attention. Our overarching position is

that the Omnibus data package represents a meaningful step towards a more risk-based, coherent, and innovation-compatible data protection framework, but that the value of many amendments will depend on harmonised guidance, consistent supervisory application, and close EDPB involvement to prevent fragmentation and to ensure that simplification does not translate into reduced protection in practice.

Executive Summary

The European AI Forum broadly welcomes the Digital Omnibus on Data as a meaningful step towards a more coherent, proportionate, and innovation-compatible data protection framework. For European AI startups and scale-ups, the current framework imposes structural compliance burdens that are poorly calibrated to actual risk: legal uncertainty about the scope of personal data, the absence of a clear legal basis for AI training, fragmented national requirements for impact assessments, and consent mechanisms that generate administrative overhead without commensurate privacy benefit. The EAIF's detailed assessment follows: where the package delivers, where its impact will depend on implementation, and where further legislative action is required.

I. Legal basis for AI development: The clarification in Art. 88c that AI development may rely on legitimate interests under Art. 6(1)(f) is essential and overdue. Legal uncertainty on this point has been a direct factor in decisions by European AI companies to train models outside the EU, without any corresponding benefit to data subjects. The provision does not create an AI exception to the GDPR: the balancing test remains mandatory, safeguards are reinforced, and the right to object is preserved. The EAIF calls on the Commission and the EDPB to frame this provision explicitly as a measure that brings legal certainty to established practice, and to publish harmonised guidance on the balancing methodology before the provision enters into force.

II. Sensitive data in AI training: The exemptions in Art. 9(2)(k) and (l) acknowledge operational realities that the current framework does not accommodate. The safeguard design requires revision: the obligation to detect and remove special category data creates a structural paradox, as detection itself constitutes processing. The EAIF calls on the co-legislators to replace prescriptive filtering mandates with a risk-based framework centred on demonstrated technical safeguards, including adversarial testing and privacy-enhancing technologies. The biometric verification exemption must be interpreted as categorically excluding remote identification in publicly accessible spaces, and the co-legislators should ensure both exemptions are accompanied by binding interpretive guidance requiring strict supervisory application.

III. Scope of personal data and pseudonymisation: The entity-relative definition in Art. 4(1) and the implementing act framework in Art. 41a together address one of the most consequential sources of overcompliance in European AI development. To deliver their

full value, Art. 41a must be extended to establish clear criteria for when data qualifies as genuinely anonymous and therefore falls outside GDPR scope entirely, and must explicitly recognise organisational and contractual safeguards alongside technical measures. Harmonised EDPB guidance must be in place before divergent national practice becomes entrenched.

IV. Harmonisation of impact assessments and breach notification: The harmonised DPIA framework in Art. 35 and the revised breach notification threshold in Art. 33 represent the most operationally significant simplification measures in the package for AI startups operating across multiple jurisdictions. Their value depends on two conditions: the DPIA framework must include an explicit grandfathering clause for existing processing operations, and must function as a ceiling on national obligations rather than a floor that Member States may supplement. The meaning of "likely high risk" under the revised Art. 33 threshold must be defined through harmonised supervisory guidance before the amendment takes effect.

V. Consent architecture: The EAIF supports the integration of ePrivacy rules into the GDPR framework under Art. 88a and the introduction of automated consent signals under Art. 88b. Two conditions are essential: the first-party analytics exemption must extend to third-party processors acting on behalf of controllers, and the browser-level obligation in Art. 88b must be restructured as a pathway rather than a mandatory requirement, with the standardisation process as the operative compliance mechanism. The consent standard under Art. 88a must remain fully equivalent to the GDPR standard and must not be read as introducing a lower threshold through the vehicle of regulatory integration.

Detailed feedback on the proposal

Art. 88c GDPR (New): Legitimate Interests as Legal Basis for AI Development

The new Article 88c clarifies that the development and operation of AI systems may rely on the legitimate interests legal basis under Article 6(1)(f), subject to the standard balancing test, enhanced technical and organisational safeguards, data minimisation requirements, and a right to object that cannot be overridden on grounds of the controller's commercial or operational interests. This provision does not create a special regime or a derogation from GDPR principles. It is a contextual clarification that confirms the applicability of an existing legal basis to AI-related processing under the same conditions that apply to all other complex processing operations.

The clarification is essential. Legal uncertainty about whether AI development can constitute a "legitimate interest" sufficient to prevail in the balancing test has been one of the most consequential practical barriers to AI development in Europe. It has driven European AI companies to train models in jurisdictions with clearer legal frameworks, to structure operations so that training occurs outside the EU, or to delay development altogether while legal positions are evaluated. None of these outcomes benefits European data subjects: they reduce European competitiveness without delivering any corresponding privacy protection, because the same personal data is processed abroad under frameworks with fewer formal safeguards than the GDPR.

The principal political risk is mischaracterisation of this provision as an "AI exception" to the GDPR. It is not. The balancing test remains mandatory on a case-by-case basis, individual safeguards are reinforced, and the right to object that cannot be overridden on grounds of the controller's commercial or operational interests preserves meaningful individual control. The Commission and EDPB should frame this provision consistently and explicitly as a measure that brings legal certainty to established practice, not as an expansion of permissible processing. Harmonised EDPB guidance on the balancing methodology for AI-related processing, developed as a matter of priority and published before the provision enters into force, is essential to prevent divergent national interpretations from recreating the uncertainty the article is designed to resolve.

Art. 9(2)(k) & (l) GDPR: Sensitive Data in AI Development and Biometric Verification

The new exemptions under Article 9(2) address two distinct challenges: the incidental presence of special category data in large-scale AI datasets, and biometric data processing for user-controlled identity verification. Both exemptions are accompanied by safeguard requirements, including technical and organisational measures to prevent

or minimise collection, obligations to remove or shield sensitive data where identified, and strict conditions governing the use of sensitive data in outputs.

The AI development exemption under Article 9(2)(k) acknowledges operational reality. In large, complex datasets used for AI training, the complete exclusion of special category data through pre-processing is technically difficult and, in some cases, functionally impossible without undermining dataset representativeness. Providing a structured legal framework for these situations, subject to proactive mitigation obligations, is preferable to leaving organisations in a position where any incidental presence of sensitive data triggers formal non-compliance regardless of actual risk.

The safeguard design in the current draft contains a structural paradox that the co-legislators must address: to comply with the obligation to detect and remove special category data from training datasets, controllers must first actively process that data. Rigid filtering mandates therefore require controllers to engage in exactly the processing they are simultaneously obliged to minimise, and strict output filters risk degrading the performance of AI systems designed to provide accurate, descriptive information about the world. The focus should shift from compliance through mandatory filtering to compliance through appropriate risk mitigation. Demonstrated safeguards such as adversarial testing, differential privacy techniques, and privacy-enhancing technologies provide a more proportionate and technically sound framework than prescriptive filtering obligations. The scope of the exemption should nonetheless be defined narrowly, with explicit limits on the types and contexts of permissible sensitive data processing, to prevent it from expanding into a general licence for special category processing in AI contexts.

The biometric verification exemption under Article 9(2)(l) raises more fundamental rights concerns. Biometric identifiers are inherently unique, largely immutable, and carry heightened risks of misuse, discrimination, and irreversible harm in the event of compromise. The "sole control of the data subject" condition is the central legitimising criterion, and it must be interpreted as categorically excluding remote biometric identification in publicly accessible spaces. The co-legislators should ensure that both exemptions are accompanied by binding interpretive guidance requiring strict application by supervisory authorities, and the Commission should resist any pressure to extend their scope through delegated acts in ways that would normalise broad reliance on special category data in AI systems.

Art. 4(1) GDPR: Entity-Relative Definition of Personal Data

The proposed clarification of Article 4(1) introduces an important and overdue refinement: identifiability becomes entity-relative, meaning that data qualifies as personal only for entities that have means "reasonably likely to be used" to identify the

individual. This codifies the CJEU's approach in Case C-413/23 P and replaces a framework in which the theoretical identifiability of data by any actor in a complex processing chain could trigger full GDPR obligations for entities with no realistic capacity for re-identification.

The practical significance for AI development, pseudonymisation-based data sharing, and cross-sectoral analytics is substantial. Under the prior framework, legal uncertainty about the scope of "personal data" drove overcompliance that imposed disproportionate costs on startups and SMEs, which typically lack the legal resources to navigate interpretive ambiguity in high-stakes regulatory environments. Overcautious data governance has been a direct factor in decisions by European AI companies to conduct training operations outside the EU, where legal uncertainty is lower. By anchoring identifiability to the specific capabilities of the relevant entity, the amendment creates a more functional and proportionate basis for responsible data reuse.

To maximise the amendment's practical effect, Recital 27 should be further clarified in two respects. First, it should explicitly confirm that purely theoretical or speculative re-identification scenarios, which are not plausible given the technical feasibility and cost of identification in real-world conditions, do not constitute "means reasonably likely to be used." Second, in line with the CJEU's ruling in the SRB case, it should make clear that the mere technical ability to single out an individual within a dataset does not render that individual identifiable where there are no actionable means to link that distinction to the individual in practice.

Without harmonised EDPB guidance specifying these criteria before divergent national practice becomes entrenched, the amendment risks recreating at enforcement level the very uncertainty it is designed to remove. Such guidance is a precondition for the amendment's effectiveness, not an optional complement to it.

Art. 41a GDPR (New): Commission Criteria for Pseudonymisation Risk Assessment

The new Article 41a gives operational effect to the entity-relative identifiability framework by empowering the Commission to adopt implementing acts specifying criteria for determining when pseudonymised data falls outside the GDPR's scope for particular entities. A harmonised EU-level risk assessment methodology would reduce fragmentation, provide controllers with a predictable compliance pathway, and enable more consistent treatment of pseudonymised data in cross-border data sharing arrangements.

The provision targets a genuine need. In complex data ecosystems, including AI training pipelines, federated analytics, and research data infrastructure, the absence of objective re-identification risk criteria has forced organisations into either overly cautious

approaches that limit legitimate data use or legally uncertain territory where legally sound compliance is impossible to achieve.

The EAIF recommends two substantive additions to Article 41a. First, the article should establish clear criteria for determining when data qualifies as genuinely anonymous and therefore falls outside GDPR scope entirely. The current framework leaves this threshold undefined, which in practice drives overcompliance: organisations treat data as personal even where re-identification is technically implausible for any realistic actor. Objective anonymisation criteria, developed with close EDPB involvement, would provide the legal certainty that the entity-relative framework in Article 4(1) is designed to deliver. Second, the article should explicitly recognise that the "means and criteria" for assessing re-identification risk are not limited to technical measures alone. Organisational and contractual safeguards, including contractual restrictions prohibiting recipients from re-identifying individuals or sharing data onward, play an equally important role in practice. Recognising the combined role of technical, organisational, and contractual safeguards would better align the legal assessment with operational realities.

The legitimacy of this mechanism depends entirely on how it is designed and maintained. Criteria for when data effectively falls outside GDPR scope must be technically robust, subject to regular review as capabilities evolve, and developed with close EDPB involvement rather than as a predominantly Commission-driven process. Robust parliamentary scrutiny of delegated acts in this area is equally essential to ensure that the mechanism does not erode fundamental rights protections under the cover of technical standardisation.

Art. 5(1)(b) GDPR: Purpose Limitation and Research-Compatible Further Processing

The amendment to Article 5(1)(b) confirms that further processing for archiving in the public interest, scientific or historical research, or statistical purposes is compatible with the initial collection purpose, independently of the conditions in Article 6(4). This is primarily a clarificatory provision. The compatibility of research-related further processing was already established in principle under the GDPR's existing framework, but persistent interpretive divergence across Member States has in practice generated significant compliance friction in data-driven research and AI development contexts.

The practical value lies in enabling responsible secondary data use without requiring data controllers to re-examine each further processing step under the Article 6(4) balancing test, provided that the Article 89(1) safeguards are met. This matters especially for AI development workflows where iterative model improvement, bias testing, and performance monitoring involve processing that was not fully anticipated at the point of initial data collection.

The amendment should explicitly confirm that it does not extend to commercial AI training conducted without a public interest rationale. Without this clarification, there is a risk that the provision is read as a general licence for secondary use of personal data in AI training, which would be both legally incorrect and politically damaging in parliamentary debate. Where commercial AI development requires a legal basis for personal data processing, Article 88c provides the appropriate framework, subject to its balancing test and safeguard requirements. Drawing this boundary explicitly pre-empts that misreading and reinforces the safeguard-based character of the provision.

The Commission should issue guidance requiring supervisory authorities to apply Article 89 safeguards as substantive requirements rather than procedural formalities, and should clarify the interaction between the amended purpose limitation principle and specific AI Act data governance obligations.

Art. 35 GDPR: Harmonised DPIA Framework

The proposed amendments to Article 35 create an EU-level harmonisation mechanism for Data Protection Impact Assessments. The EDPB would develop Union-wide lists of processing operations requiring or not requiring a DPIA, along with a common template and methodology, subsequently adopted through Commission implementing acts. This replaces the current system in which national supervisory authorities maintain separate and divergent lists.

For organisations operating across borders, this reform has substantial practical value. The fragmentation of DPIA requirements has created genuine compliance complexity: identical AI processing activities may require a DPIA in some Member States but not in others, and the divergence has been particularly acute for AI systems that are deployed across multiple jurisdictions simultaneously. A harmonised Union-level framework removes this uncertainty and provides a consistent baseline for risk assessment practice across the internal market.

Two conditions are critical to the quality of outcomes. First, the development of lists and methodology must be genuinely evidence-based and responsive to technological change. Lists that reflect only the processing landscape at the time of adoption will rapidly become outdated as AI capabilities evolve, and outdated lists generate false compliance confidence. Second, the implementing act framework must include an explicit grandfathering clause confirming that new DPIA obligations apply only to new processing operations. Existing processing that has been operating in compliance with current national guidance cannot logically be subjected to retroactive assessment: it is structurally impossible to conduct a DPIA "prior to processing," as Article 35 requires, for operations that have been running for years. Without this clause, companies face legal

uncertainty about the status of established services, and the harmonisation exercise creates new compliance risks rather than reducing them.

The relationship between Union-level implementing acts and residual national supervisory discretion should be clearly defined to prevent the harmonised framework from functioning as a floor that Member States can supplement with additional national obligations, which would recreate the very fragmentation the amendment is designed to remove. The EDPB's harmonised list should reflect the existing consensus of national supervisory authorities rather than introducing new, expansive requirements that aggregate all national lists into a more burdensome combined instrument.

Art. 33 GDPR: Breach Notification Threshold and Harmonised Reporting

The proposed amendment raises the supervisory authority notification threshold from "risk" to "likely high risk" to the rights and freedoms of natural persons, extends the notification deadline to 96 hours, and introduces a harmonised EU-level reporting mechanism with a single entry point and a common notification template.

The threshold adjustment addresses a well-documented operational problem: the volume of low-risk breach notifications submitted under the current framework diverts supervisory resources from incidents with genuinely serious impact and imposes compliance costs on organisations for reporting events that present negligible risk to individuals. For European AI startups operating with lean security teams, the current framework creates a reporting burden that is poorly calibrated to actual risk levels: minor incidents in development environments trigger the same formal notification obligations as serious data exposures. Aligning the supervisory notification threshold with the Article 34 data subject notification standard introduces greater internal GDPR coherence and ensures that supervisory attention remains proportionate to actual risk.

Extending the notification deadline to 96 hours reflects operational reality for complex incidents involving distributed systems or third-party processors. Notifications that are complete and accurate deliver greater supervisory value than notifications that are fast but fragmentary. The harmonised reporting mechanism and common template have the potential to substantially reduce fragmentation and cross-border compliance complexity, benefits that are particularly significant for startups and scale-ups operating across multiple jurisdictions.

To prevent the threshold change from reducing investment in detection infrastructure, robust internal incident response obligations and consistent supervisory guidance on the meaning of "likely high risk" are necessary counterparts. The Commission should also clarify that the reduced formal reporting burden does not in any way diminish controllers' obligations to document, assess, and remediate all security incidents internally, regardless of whether they reach the notification threshold.

Art. 12 GDPR: Addressing Abusive Data Subject Access Requests

The amendment to Article 12 clarifies that controllers may refuse or charge a reasonable fee for data subject access requests that are manifestly unfounded, excessive, or abusive, including where access rights are used for purposes unrelated to data protection. The burden of demonstrating that the relevant threshold is met remains with the controller.

For AI startups and scale-ups, the absence of harmonised standards for assessing request legitimacy creates a structural disadvantage: without in-house legal capacity, companies face significant uncertainty in distinguishing genuine from unfounded requests, and bear the full cost of compliance with requests that serve no substantive privacy interest. The amendment's practical value depends on narrow, consistent application. Overly permissive interpretations of "abusive" could deter legitimate exercise of rights by data subjects who lack the resources to challenge controller refusals, undermining the right of access as a tool of individual control. EDPB guidance defining objective parameters for demonstrable abuse, and establishing clear procedural safeguards for refused requests including timely review mechanisms, should be published before the provision enters into force.

Art. 13 GDPR: Proportionate Information Obligations

The amendments to Article 13 introduce targeted simplifications of information obligations in low-risk contexts: where personal data are processed within a clearly circumscribed relationship and where the data subject can reasonably be assumed to already possess the relevant information, certain mandatory disclosure elements may be reduced. A parallel clarification addresses the disproportionate effort exemption for scientific research involving large-scale secondary data use.

The proportionality rationale is sound. Repeated, highly formalised information notices in stable, low-intensity processing relationships often provide limited additional value to data subjects while generating compliance costs that fall disproportionately on smaller organisations. The amendment is consistent with the GDPR's existing risk-based approach and does not alter the fundamental transparency obligations that apply in higher-risk processing contexts.

The critical implementation risk is definitional. The terms "non-data-intensive" activity and "circumscribed relationship" are not defined in the proposal, and without precise criteria they will be applied inconsistently or opportunistically across Member States, recreating the fragmentation the amendment is intended to reduce. The EAIF recommends that the Commission adopt a definition anchored to the AI Act's risk classification framework: where a processing activity underlies an AI system that is not classified as high-risk under the AI Act, it should presumptively qualify for the simplified

information regime, subject to EDPB guidance on sector-specific exceptions. This approach links two regulatory frameworks coherently, provides a workable threshold for controllers, and enables supervisory authorities to calibrate their oversight proportionately.

The research-related clarification addresses a genuine operational constraint. In large-scale secondary data research, individually contacting data subjects may be technically impracticable or disproportionately costly relative to the research benefit. Linking the exemption to Article 89 safeguards, including public availability of relevant information, maintains the transparency principle while acknowledging the realities of contemporary AI research pipelines. For AI startups offering API-based services, which typically interact with large numbers of data subjects through standardised interfaces rather than individual relationships, the simplified information regime is particularly relevant: the current framework requires repeated, highly formalised disclosures in contexts where data subjects already understand the nature of the service and repeated notices deliver no meaningful additional protection.

Art. 22 GDPR: Automated Decision-Making and Contractual Necessity

The amendment clarifies that a decision producing legal or similarly significant effects may be based solely on automated processing where it is necessary for entering into or performing a contract, without requiring the controller to demonstrate that the decision could not theoretically have been taken by human means. This addresses interpretive divergence across Member States about the scope of the "contractual necessity" exception under Article 22(2)(a).

The justification for this clarification rests not primarily on efficiency grounds, but on quality grounds: automated decisions in well-defined, data-rich contexts can be demonstrably more consistent, more objectively calibrated, and less susceptible to individual bias than equivalent human decisions. Algorithmic credit assessment, for example, applied to clearly defined criteria and subject to auditability requirements, can produce demonstrably fairer outcomes for applicants than discretionary human review, particularly in high-volume, rule-based contexts where individual reviewer bias and inconsistency are documented risks.

The central safeguard concern is that a broad reading of "contractual necessity" could progressively displace the human oversight that Article 22 was designed to preserve. Necessity must remain a genuine criterion, assessed by reference to whether automation serves a legitimate and demonstrable function. Supervisory authorities should apply strict scrutiny to contractual necessity claims, and the rights to obtain human intervention, to contest decisions, and to receive meaningful explanations must be enforced in full. The EAIF recommends that the Commission publish guidance on

what constitutes demonstrable contractual necessity in specific AI deployment contexts, to prevent the provision from being read as a general permission to automate decisions that happen to be associated with a contractual relationship. This is of particular relevance to AI startups developing credit scoring, insurance underwriting, and recruitment screening tools, where automated decisions are both the core product and the primary contractual deliverable, and where the boundary between genuine contractual necessity and commercial preference for automation is most likely to be contested.

Art. 88a GDPR (New): Cookies and Terminal Equipment Access

The new Article 88a integrates rules on accessing or storing data in terminal equipment, currently governed by the ePrivacy Directive, directly into the GDPR framework. It maintains the consent requirement in principle while introducing clearly defined exceptions and procedural safeguards, including single-click refusal mechanisms and a minimum period before re-requesting consent.

The regulatory coherence argument for this consolidation is strong. The dual GDPR-ePrivacy regime has long generated interpretive complexity, inconsistent national implementation, and enforcement confusion about which rules apply where personal data are involved. A unified legal framework with a single enforcement architecture reduces fragmentation and improves predictability for controllers operating across multiple jurisdictions. The standardised refusal mechanisms and cooling-off period address consent fatigue, which is one of the most widely documented practical failures of the current framework. The proliferation of intrusive consent banners has produced a culture of reflexive interaction rather than informed choice. For AI developers and data-driven service providers, this matters directly: meaningful consent is the foundation of a legitimate data ecosystem, and restoring its substance benefits the entire data economy, not merely individual users.

The EAIF recommends two targeted clarifications. First, the exemption for first-party analytics should explicitly extend to processing carried out by specialist third-party processors or controllers acting on behalf of the service provider, not only to processing conducted solely by the controller itself. Restricting the exemption to in-house analytics is more stringent than existing national guidance and does not reflect how analytics services are organised across the European digital economy. Second, the consent standard under Article 88a must remain fully equivalent to the GDPR standard. This provision must not be read as introducing a lower consent threshold for tracking technologies through the vehicle of regulatory integration.

The six-month consent renewal period should be replaced by a risk-based standard that reflects the diversity of digital services and technical environments. A uniform fixed

period cannot be consistently implemented across authenticated and unauthenticated services, across browser-based and app-based environments, or where users delete locally stored privacy preferences. A proportionality standard calibrated to the sensitivity of the processing activity would better serve both user protection and operational predictability.

Art. 88b GDPR (New): Automated Consent Signals

The new Article 88b requires controllers to respect machine-readable, automated signals expressing data subjects' consent choices or objections, where such signals conform to standards developed at EU level. It further requires certain providers, including web browsers subject to defined scope limitations, to enable users to express preferences through technical means. Compliance with harmonised standards creates a presumption of conformity.

This provision addresses a structural shortcoming in the current consent model. Banner-based consent interactions are widely recognised as failing to produce informed, meaningful choices. The combination of dark patterns, decision fatigue, and the asymmetry between the ease of accepting and the effort of refusing consent has produced an environment in which formal consent is routinely obtained without genuine user understanding or reflective preference expression. Technical standards enabling persistent, device-level preference signals have the potential to restore substance to consent as a legal basis.

The EAIF supports the underlying objective of Article 88b but recommends that the co-legislators restructure its architecture in two respects. First, the browser-level obligation should be framed as a pathway to enable standardised signals rather than as a mandatory requirement in its current form. Browser-level consent choices cannot operate reliably in cross-device environments: a user who withdraws consent within an authenticated mobile application may continue to have that same service receive a "consent granted" signal from a browser with no technical capacity to reconcile the updated preference. This creates compliance conflicts that disadvantage controllers and undermine, rather than strengthen, the persistence of user privacy preferences. Second, the standardisation process should be made the operative mechanism of the article: standards developed through an open, multi-stakeholder process with genuine civil society, technical, and industry participation should define what constitutes a valid automated signal, and conformity with those standards should create the presumption of compliance. This approach is technically more realistic, allows the framework to evolve with the technology, and avoids binding controllers to requirements whose feasibility depends on technical infrastructure that is still under development.

Default settings in browsers and devices must be structured so that automated signals genuinely reflect informed user preferences rather than commercially motivated defaults. The fundamental rights value of Article 88b depends entirely on whether its technical implementation preserves the individual autonomy it is designed to enable.

The European AI Forum remains available to elaborate on any of the positions set out in this statement and welcomes further engagement with the Commission, the co-legislators, and the EDPB in the course of the legislative process.

About the European AI Forum EAIF:

The European AI Forum (EAIF) is the umbrella organisation of 13 national European AI associations and clusters. Combined, we represent over 3,000 members, making us the largest European AI organisation. We represent members including SMEs, corporate members, organisations and individual experts in AI.

Contact:

info@eaiforum.org