

Brussels, 13th March 2026

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

Digital Omnibus on AI¹

Executive Summary

The European AI Forum broadly welcomes the Digital Omnibus on AI as a necessary corrective to an implementation framework that has imposed disproportionate burdens on European startups, SMEs, and Small Mid-Cap Companies. The Commission's projected savings of up to €5 billion² by 2029 reflect the genuine scale of the problem. 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. Implementation timeline: The Commission-trigger mechanism in Art. 113(3)(d) introduces new uncertainty: companies cannot plan around a Commission decision whose timing is entirely within the Commission's discretion. The EAIF calls on the co-legislators to write fixed dates directly into the operative text: 2 December 2027 for Annex III systems and 2 August 2028 for Annex I. These dates must be treated as firm outer limits, backed by a formal Commission obligation to publish the necessary harmonised standards before those dates. Where that obligation cannot be met, the dates should shift accordingly through an automatic and transparent adjustment mechanism tied to the publication of standards, not through discretionary Commission action.

II. Proportionality along the supply chain: Extending proportionality measures to Small Mid-Cap Companies is one of the most significant structural improvements in the package. The EAIF urges the co-legislators to complement it by addressing liability allocation within AI value chains: compliance obligations must not be contractually displaced onto smaller providers who lack the capacity and control to discharge them.

III. Regulatory sandboxes: The Union-level sandbox with priority access for SMEs and startups is a welcome development whose value depends entirely on accessibility and on what participation delivers. The EAIF calls for proportionate entry requirements, binding admission timelines, and a formal presumption of conformity upon successful sandbox exit.

¹ Regulation amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as regards the simplification of the implementation of harmonised rules on artificial intelligence. COM(2025)836.

² https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2718.

IV. Avoiding duplication: Where a GDPR Data Protection Impact Assessment already fulfils the substance of a Fundamental Rights Impact Assessment, it should satisfy the Art. 27 obligation in whole or in part. More broadly, the EAIF calls for a "fill in once" approach to regulatory risk assessments: the foundational technical information required across DPIAs, Transfer Impact Assessments, Legitimate Interest Assessments, and FRIAs is substantially the same, and organisations should not be required to populate it separately for each framework. Where established sectoral authorities hold the expertise and mandate to assess AI systems in their domain, conformity assessment should be conducted through those frameworks rather than in parallel to them.

V. The integrity of the framework: The EAIF does not support amendments that remove substantive obligations without demonstrated justification. The AI literacy obligation should be retained and clarified, not deleted. The AI Office requires a formal mediation and harmonisation role for cases where national competent authorities arrive at divergent interpretations of the same provision, with a mandatory consultation mechanism, a defined timeline, and an obligation to issue a binding or advisory opinion. Without such a mechanism, enforcement fragmentation across Member States will generate precisely the legal uncertainty this package seeks to address.

Detailed feedback on the proposal

Europe's ability to develop and deploy artificial intelligence at scale is a strategic question, not only an economic one. The United States and China are investing massively in their AI ecosystems, and the companies that will define the next generation of AI infrastructure, foundation models, and applications are being built today. Whether they are built in Europe or whether European talent, capital, and data flow elsewhere depends significantly on whether the regulatory environment enables European AI companies to move at speed. European AI sovereignty is not secured by regulation alone; it is secured by the existence of competitive, scaling European AI companies that can develop sovereign capabilities, maintain European data infrastructures, and anchor AI value chains within the single market. A regulatory framework that systematically disadvantages European startups and SMEs relative to their international competitors undermines both.

The Digital Omnibus simplification package is a necessary corrective to a framework that has imposed disproportionate burdens on European AI startups and SMEs since its entry into force: complex conformity assessment requirements, unclear interfaces with sectoral legislation, ambiguous data governance rules, insufficient transitional arrangements, and enforcement regimes calibrated for large multinational enterprises. These burdens are not abstract. A representative of a mid-sized industrial AI company describes investing "enormous resources in personnel and external consultancy to ensure compliance, without any of this generating real value for the product." A legal technology SME reports that the volume of applicable regulatory material running to approximately one thousand pages, already exceeds the capacity of a small organisation to study systematically, let alone implement.³ These cases are representative, not exceptional. The Commission's projected reduction in annual administrative costs of up to €1 billion, cumulating to approximately €5 billion by 2029⁴, reflects the scale of the structural problem the Omnibus is designed to address. Several of the amendments proposed in this package directly respond to positions the EAIF has advocated for: the removal of the pre-registration obligation for non-high-risk systems under Art. 49(2), the replacement of the blanket Art. 4 literacy obligation with a proportionate policy mandate, the extension of sandbox frameworks, and the introduction of conditional commencement mechanisms for high-risk obligations. These are meaningful steps, and the EAIF acknowledges them as such. This paper sets out the EAIF's further detailed assessment of each proposed amendment of the Digital Omnibus Proposal on AI.

³ Hacker, P., Kilian, R. & Costas, J. (2025) Simplifying European AI Regulation — An Evidence-based Study. Gütersloh: Bertelsmann Stiftung. Interview findings reported anonymously in accordance with research ethics protocols.

⁴ https://ec.europa.eu/commission/presscorner/detail/en/jp_25_2718.

Art. 1(2)(g) AI Act – Extension of Innovation-Support Clause to Small Mid-Cap Companies

The original innovation-support clause focused only on SMEs, leaving out a critical group: companies that have just grown beyond the SME threshold but continue to face the same challenges in compliance, capital access, and scaling. Extending the clause to SMCs closes that gap. The so-called "cliff effect" is a genuine structural problem. Companies that marginally exceed the SME threshold lose their regulatory flexibilities overnight, and for European AI companies this can happen rapidly: a single funding round or period of rapid hiring can push a company over the line before its compliance infrastructure has had time to catch up. Recognising SMCs as a distinct category with proportionate treatment aligns the rules with how companies actually grow, and the EAIF urges the legislators to preserve this extension in full. A separate issue the Parliament has rightly identified concerns the distribution of obligations within AI value chains. The proportionality framework is designed to protect smaller providers from disproportionate compliance burdens, but it cannot achieve that objective if larger deployers remain free to contractually reassign those burdens downward, irrespective of where actual control over the AI system lies. A provider that has no meaningful influence over how its system is deployed cannot reasonably be held to compliance obligations that presuppose such control. The EAIF strongly supports parliamentary amendments addressing this point. Proportionality must be structurally protected along the full supply chain, not merely available in principle while being systematically circumvented in practice.

Art. 3(14a) & (14b) – Legal Definitions of SMEs and Small Mid-Cap Companies

Putting clear definitions of SME and SMC directly into the AI Act creates legal certainty where there was none before. Both definitions now have a single, authoritative source, which means competent authorities and companies across all Member States work from the same baseline. This prevents the divergent national interpretations that have plagued earlier frameworks. Recognising SMCs as a distinct legal category, rather than lumping them in with large enterprises, has direct practical consequences. It determines which companies can access simplified procedures, regulatory flexibilities, and targeted support. Without this clarity, those commitments risk being hollow in practice. For fast-growing AI companies, this also helps with planning. A company that knows exactly when SMC status kicks in and what obligations follow can prepare in advance rather than scrambling to catch up. That predictability reduces the compliance cost that currently hits scaling companies hardest and makes the EU a more attractive place to build.

Art. 3 AI Act (Parliamentary Amendment) – Extending the Definition to AI Agents

The European Parliament's JURI committee has proposed extending the AI Act's core definition of an AI system to explicitly cover systems that take autonomous actions in the world. The underlying concern is legitimate: as AI systems become more capable of executing multi-step tasks and acting without continuous human instruction, a definition anchored only in output generation risks creating gaps in the regulatory perimeter as the technology matures.

The EAIF supports addressing this gap, but urges the co-legislators to do so in a technology-neutral way. Anchoring the definition to a specific system architecture, such as "AI agents" as a named category, risks becoming obsolete as the technology evolves and generates interpretive uncertainty about where the boundary lies. The more durable approach is to ensure that the existing definition of an AI system is broad enough to capture autonomous action-taking as a form of output, without introducing a separate named category that will require constant revision. Future-proofing the framework means regulating the capability, not the architecture.

Art. 2(2) AI Act – Application Regime for Product-Regulated AI Systems

AI systems built into regulated products (e.g. medical devices, vehicles, or industrial machinery) sit at the intersection of the AI Act and sectoral product law. Until now, it was unclear which AI Act rules applied to them and to what extent. This amendment answers that question. It specifies exactly which AI Act provisions apply: Article 60a, and Articles 111 and 112. Everything else is governed by the sectoral framework. This matters because the uncertainty was not theoretical. Companies building AI for regulated product markets had real difficulty knowing whether AI Act conformity obligations applied on top of their existing sectoral assessments. The amendment makes clear that they do, but in a defined and complementary way, not as a second full compliance track. The sandbox rule is also sensibly scoped: Art. 57 only applies where sectoral legislation has already integrated AI-specific requirements. This prevents companies from using sandbox access as a workaround and keeps the instrument focused on situations where it adds genuine value. For startups building AI for regulated product markets, this amendment removes a persistent source of uncertainty that has held back investment decisions and delayed product roadmaps.

Art. 4 AI Act – AI Literacy

The original Art. 4 imposed an open-ended AI literacy obligation on all providers and deployers, including micro-enterprises and startups, without any differentiation by size,

sector, or context. The amendment shifts that responsibility to the Commission and Member States, who are better placed to develop accessible training resources at scale, and moves from a mandatory obligation to a voluntary approach that removes the compliance burden on companies while keeping the policy objective intact. The risk of uneven uptake is real but manageable and far outweighed by the proportionality gain. For this to work, the Commission and Member States must deliver substantive, practice-oriented initiatives rather than formal outreach. Binding measures should remain available as a backstop if voluntary uptake proves systematically insufficient.

The EAIF does not support the complete removal of the AI literacy obligation, as proposed by some amendments before Parliament. Independent empirical research confirms that the more defensible approach is to retain the obligation while making it workable: clarifying its scope and personal application, issuing sector-specific guidance, and differentiating expectations by company size and risk context. In particular, an exemption for administrative staff of providers and deployers would ensure proportionality without undermining the obligation's substantive purpose. Removing the obligation altogether risks being read as indifference to responsible AI use, which is neither accurate nor strategically sound for the European AI sector.

One further reform the EAIF calls for: where companies already hold valid GDPR certifications covering processes that also fall under AI Act documentation requirements, those certifications should count. Requiring companies to certify the same processes twice under different labels generates cost without generating any additional safety or transparency benefit. Beyond the obligations on providers and deployers, the EAIF also calls for Art. 4 to be read as the foundation for a broader societal ambition: genuine AI literacy cannot be achieved by targeting only the companies that build and deploy AI systems. The Commission and Member States should promote measures aimed at the general public, and recommendations from the AI Board setting out practical frameworks for public AI literacy programmes would provide a useful complement to the obligations on providers and deployers without adding new compliance burdens on business.

Art. 4a AI Act (New) – Processing of Special Categories of Data for Bias Detection

Art. 4a is trying to close a conflict between the AI Act and the GDPR. The AI Act requires providers to detect and correct bias in training data, but doing so often means processing sensitive personal data like racial or ethnic origin, which the GDPR tightly restricts. Art. 4a creates a safeguarded exception: providers may process such data, but only for bias detection and only under strict conditions. Without the ability to use relevant data for bias detection, meaningful discrimination prevention is not achievable

in practice and the risk of undetected bias causing harm to real people outweighs the risk of controlled, safeguarded processing for this specific purpose.

The EAIF supports the "necessary" standard for this exception rather than a more restrictive formulation. A "strictly necessary" threshold risks making bias testing unworkable in practice, requiring companies to demonstrate at each step that no alternative existed, a burden that would deter precisely the proactive bias management that the provision is designed to encourage. "Necessary" remains a genuine proportionality standard; it is not a blank licence. The safeguards are robust: data must be minimised, deleted once the bias correction is complete, and kept strictly within the organisation.

Notably, Art. 4a(2) extends the exception beyond high-risk AI systems, recognising that discriminatory outputs can emerge in any AI application, not only in formally high-risk ones. The EAIF explicitly supports this extension. Confining bias-detection rights to high-risk systems would create a regulatory blind spot: systems that fall below the high-risk threshold can still produce outputs with significant discriminatory effects on real people. The safeguards attached to Art. 4a(2) ensure that the extension does not become a licence for unrestricted sensitive data processing; it remains a narrow, purposive permission tied to discrimination prevention. The EAIF calls on the co-legislators to retain and strengthen this provision in the final text.

Previously, the relevant rules were scattered across Art. 10(5) and other provisions. Consolidating them into a single, self-contained article makes compliance significantly more accessible, especially for smaller companies without dedicated legal teams who previously had to cross-reference multiple frameworks just to understand what was permitted. This is one of the most practically important amendments in the package.

Art. 6(4) / Art. 49(2) – Streamlined Classification for Borderline High-Risk Systems

Under the previous rules, a company that assessed its AI system as non-high-risk was nonetheless required to register it in the EU database where the system could potentially fall within the scope of Annex III — that is, where it belonged to a category of systems that the Act treats as presumptively high-risk, even if the provider had concluded, after self-assessment, that the specific system did not meet that threshold. This created administrative overhead with no regulatory purpose: the registration obligation applied precisely in cases where the assessment had already determined that no high-risk obligations were triggered.

The amendment removes that obligation. Accountability is nonetheless preserved: companies must continue to document their risk assessment and make it available to

competent authorities on request, meaning that documentation now serves a genuine supervisory function rather than operating as a default administrative exercise triggered by uncertainty about classification. The risk of under-classification is managed through that retained documentation obligation, since authorities can still demand access and challenge assessments they consider inadequate. For startups and smaller providers, the amendment removes a recurring compliance cost that bore no relationship to the actual risk posed by their systems.

The EAIF notes that the Parliament's JURI committee has taken a different view, opposing the removal of the registration obligation under Art. 49(2). The EAIF does not share that position. The registration of systems that have already been assessed as non-high-risk serves no supervisory purpose that the retained documentation obligation does not already fulfil. Accountability and the removal of administrative overhead are not in conflict here, and the EAIF urges the co-legislators to preserve the Commission's proposed amendment.

Art. 11(1) AI Act – Simplified Technical Documentation for SMEs and SMCs

SMCs can now use the simplified technical documentation format, alongside SMEs. Previously, a company that had just grown beyond the SME threshold faced a sharp jump in documentation requirements, with no transition period, regardless of whether its actual capacity had changed. This amendment removes that cliff.

The requirement for a standardised Commission template, which notified bodies must accept, is key. Without standardisation, different notified bodies across Member States could apply different interpretations of what “simplified” means, fragmenting the market and creating unpredictable compliance demands. A common template lets companies invest in the substance of compliance rather than navigating procedural inconsistency.

One implementation risk worth flagging: the template must not be so stripped down that it becomes meaningless. Notified bodies need enough information to run genuine assessments. A template that asks too little undermines the conformity process and ultimately harms the companies that depend on credible certification. The Commission should develop the template with direct input from notified bodies and SME representatives.

Art. 17(2) AI Act – Proportionate Quality Management Systems for SMEs and SMCs

Before this amendment, proportionality in quality management systems was technically implied but not stated. That left supervisory and conformity assessment bodies free to apply full large-enterprise standards to small providers, with no formal legal basis for companies to push back. The amendment makes the entitlement explicit: SMEs and

SMCs may implement a reduced set of QMS processes proportionate to their size. For startups and scale-ups this means their internal processes can be calibrated to what they actually need, rather than to a maximalist standard built for large corporations.

The Commission's commitment to guidance on which QMS elements may be simplified is welcome, but it must be delivered promptly and with real specificity. In practice, conformity assessment bodies consult guidance far more than statutory text. Generic guidance will leave companies no more certain than before. The EAIF calls for sector-specific examples and concrete benchmarks. The underlying problem is clear: where QMS requirements are disproportionate, resources that should go into product development and safety testing are consumed instead by documentation and external consultancy. Proportionality on paper is not enough, the guidance must make it real.

Art. 28 (New Para. 8) AI Act – Streamlined Designation of Notified Bodies

A persistent bottleneck in AI Act implementation has been the limited availability of notified bodies. Getting designated as a notified body under both the AI Act and existing sectoral product legislation was a costly, duplicative process. This amendment allows a single application and single assessment procedure to cover both. More notified bodies means less queuing for conformity assessment.

The new Annex XIV coding system specifies exactly which types of AI systems each notified body is qualified to assess. This makes it straightforward for companies to identify the right assessment partner from the outset, rather than discovering late in the process that a given body lacks competence for their system type.

The Commission can update the Annex XIV categories via delegated act as the technology evolves, which is the right approach given how quickly AI capabilities are developing. For startups entering new application areas, knowing which notified body is competent for their system type is not a minor convenience. It is a prerequisite for market entry planning. This provision directly reduces that uncertainty.

Art. 43(3) AI Act – Conformity Assessment for Product-Embedded AI Systems

For AI systems built into regulated products, it was previously unclear whether QMS obligations under Art. 17 applied in addition to sectoral conformity assessment. This amendment confirms they do, closing a gap that could otherwise have been used to avoid AI-specific quality requirements.

The 18-month transitional period for sectoral notified bodies to obtain formal AI Act designation is a sensible balance: it keeps assessment capacity available immediately while setting a firm deadline for formalising AI-specific expertise.

The self-certification rules are also tightened: using product law self-assessment as a route to bypass AI-specific conformity obligations is explicitly ruled out. And where a system falls under both Annex I and Annex III, the product conformity procedure takes precedence and eliminating the procedural duplication that providers of product-embedded AI have long faced when planning their compliance pathways.

Art. 50(7) AI Act – Codes of Practice for Transparency Obligations in Generative AI

The AI Office takes a stronger coordinating role in developing Codes of Practice for generative AI transparency obligations, while the Commission retains the power to adopt binding rules if voluntary measures fall short. This is the right structure: industry-led codes can respond to technical realities more precisely than prescriptive legislation, but they need a credible backstop to function.

Codes of Practice work well for transparency obligations in generative AI because the technical landscape is moving quickly. A code developed with industry can track what is actually feasible for content labelling and detection in ways that hard-coded legislation cannot. The AI Office's coordinating role ensures consistency across the Union rather than a patchwork of sector-specific arrangements.

The Commission's backstop power to adopt binding rules is necessary, but must be used with discipline. It should activate when voluntary codes demonstrably fail, not as a default preference for formal legislation. Deploying it prematurely would undermine the development of technically sophisticated voluntary standards before they have had time to mature.

Art. 56(6) AI Act – Enhanced Oversight of Codes of Practice

Publishing the Commission's evaluation of Code adequacy creates real accountability: companies and the public can see whether voluntary governance is delivering. Reputational incentives matter since Code participants have an interest in being seen to meet their commitments.

For this to work, evaluations must be substantive and not performative. Assessments that flag problems without offering actionable guidance risk discouraging the industry engagement that makes Codes of Practice function. The Commission should treat its oversight role as a collaborative process, not a gate-keeping one.

Close coordination with the AI Board is essential here. In a field where the adequacy of transparency mechanisms depends on algorithmic and infrastructural specifics, evaluations conducted without deep technical engagement will miss what matters. The Parliament has also proposed replacing certain implementing act empowerments with

obligations to issue guidance, on the basis that guidance is more flexible and faster to update as technology evolves. The EAIF supports this direction where the harmonisation objective is preserved and legal clarity is not sacrificed. For codes of practice in particular, guidance-based oversight is more likely to produce technically credible outcomes than prescriptive implementing rules drafted without sufficient operational input.

Art. 57 & 58 AI Act – AI Regulatory Sandboxes

The sandbox framework gets a substantial upgrade. From 2028, the AI Office can establish a Union-level sandbox targeting AI systems under its direct supervision, including general-purpose AI models and AI integrated into major digital platforms. SMEs, start-ups, and SMCs get priority access. This is the right call: these are the companies that most need structured regulatory engagement during development but are least equipped to navigate it alone.

At national level, sandbox programmes can now explicitly include real-world testing plans, resolving a longstanding ambiguity about how far sandbox activity could extend into actual deployment conditions. Member States are also required to cooperate on cross-border and joint sandboxes, which matters because fragmented national programmes have historically imposed inconsistent and duplicative conditions on companies operating in multiple jurisdictions. The European Parliament's JURI committee has proposed strengthening Union-level sandbox coordination further, with the aim of preventing 27 divergent national interpretations of the same instrument. The EAIF strongly supports this direction: a genuinely harmonised sandbox framework, with common entry criteria, coordinated admission processes, and mutual recognition of sandbox outcomes across Member States, is the only version of the instrument that works for companies operating across borders.

The EAIF welcomes the revised framework but flags the central implementation risk clearly: sandboxes that require extensive procedural investment to enter will systematically exclude the SMEs and start-ups they are designed to serve. The value of this instrument depends entirely on accessibility. Administrative entry requirements must be proportionate, timelines for sandbox admission decisions must be published and enforced, and the operational burden of participation must not replicate the compliance overhead companies enter sandboxes to navigate. The co-legislators should mandate that the Commission report on SME and start-up participation rates within two years of the Union-level sandbox becoming operational. One structural gap in the revised framework requires explicit correction. A successful exit from a regulatory sandbox should confer a formal presumption of conformity on the tested system. The current text treats sandbox completion as evidentiary weight only. This undervalues the investment companies make to participate and offers no regulatory certainty at the

point it matters most: market entry. A presumption of conformity would not only make sandboxes more attractive to the companies they are designed to serve, it would give the instrument the regulatory teeth it currently lacks. The EAIF calls on the co-legislators to introduce this mechanism explicitly.

Art. 60(1) & (2) AI Act – Real-World Testing for Product-Embedded High-Risk AI

Stand-alone high-risk AI systems listed in Annex III have always been able to conduct real-world pre-market testing. This amendment extends the same option to AI systems embedded in sectoral products under Annex I, Section B, covering areas like automotive and medical devices. Excluding them was never justified on principle, and was particularly damaging in sectors where realistic field conditions are essential for demonstrating conformity.

Product safety obligations are not suspended during testing. Real-world testing cannot be used as a route around existing safety frameworks, and this boundary is clear and well-drawn.

The inclusion of intended deployers in the testing process is a practically important detail. For AI systems embedded in complex products, how and where they are actually used, by whom, in what environment, and alongside what other systems, materially affects performance and safety. Involving deployers early produces better test data and more realistic conformity evidence.

Art. 60a AI Act (New) – Voluntary Real-World Testing Agreements for Safety-Critical Products

Art. 60a creates a new instrument: Commission-led voluntary testing agreements for AI systems in vehicles, aircraft, and analogous safety-critical equipment (Annex I, Section B products). These agreements allow supervised field testing outside the standard conformity assessment track, under conditions agreed jointly by providers and competent authorities.

The need is real. In automotive and aviation, validating AI-enabled functionality requires extensive real-world exposure, but the existing regulatory frameworks were designed before iterative AI testing cycles existed. Art. 60a creates a legally secure pathway for controlled testing without displacing underlying safety requirements.

Whether this instrument delivers depends on how lean the agreement process is in practice. If negotiating a voluntary testing agreement costs as much time and effort as the standard authorisation pathway, companies will not use it. The Commission should

develop standardised templates and defined timelines so that testing agreements function as actual innovation accelerators, not an additional bureaucratic layer.

The practical effect of the real-world testing framework will remain limited, however, as long as the underlying problem of data access is not addressed. Testing and validating AI systems under realistic conditions requires access to high-quality, representative datasets. For many providers, particularly in healthcare, transport, and public services, obtaining that access under legally secure and practically manageable terms is one of the most significant operational obstacles they face. The Commission should treat coordinated data access as an integral part of the testing framework, not a separate policy area. Two further gaps should be addressed explicitly: the framework should include standard contractual terms for cooperative testing between providers and deployers, establishing clear roles, liability allocation, and risk distribution. Without such templates, the uncertainty of collaborative testing arrangements will continue to deter participation, especially from smaller providers who cannot absorb the cost and delay of negotiating bespoke agreements for each testing relationship.

Art. 63(1) AI Act – Simplified QMS Regime Extended to All SMEs

The simplified QMS regime was previously available only to micro-enterprises. This amendment extends it to all SMEs, including start-ups. The previous cut-off made little practical sense: companies just above the micro-enterprise threshold often face QMS complexity far greater than their marginal increase in size would suggest, while remaining firmly in the category of Europe’s most innovation-intensive AI developers.

The safeguard excluding companies with partner or linked enterprises is well-targeted. It stops large corporate groups from routing products through formally small subsidiaries to access a regime that was designed for genuinely independent SMEs.

The Commission’s guidance on what “simplified” means in practice must be treated as a priority deliverable, and it must be specific. Generic guidance that leaves companies uncertain about what they actually need to do reproduces the problem the amendment is meant to solve. Sector-specific examples and concrete benchmarks are essential. The evidence base is clear: legal technology SMEs report that the sheer volume of applicable regulatory material, across primary legislation, delegated acts, and implementing measures, already exceeds what a small organisation can systematically study, let alone implement in full.¹ A simplified QMS regime is not a concession to companies unwilling to comply. It is the only framework under which meaningful compliance by resource-constrained innovative companies is genuinely achievable.

Art. 70 AI Act – Proportionality in Enforcement Extended to SMCs

Without this clarification, a company that had recently crossed the SME threshold, often through successful growth, not any change in how it operates, would automatically face enforcement practices calibrated for large multinationals. The financial exposure relative to actual capacity would have been severe, and in some cases existential. This amendment extends proportionality requirements to SMCs, closing that gap. This does not dilute the enforcement framework. Compliance obligations remain unchanged. What changes is that authorities must consider the financial and operational capacity of SMCs when setting penalties. That is not a loophole, it is rather good enforcement design. Fines that threaten a company's survival create incentives to conceal non-compliance rather than fix it.

The extension of this safeguard to SMCs also strengthens investor confidence in the European AI ecosystem. Founders and early-stage investors need to be able to model regulatory exposure with reasonable precision; a framework in which penalty risk increases discontinuously at the SME boundary introduces avoidable uncertainty into investment decisions. The amendment removes that distortion. The Omnibus does not address one enforcement gap that disproportionately affects the companies this framework is designed to protect. A European startup that makes a genuine, first-time compliance mistake, without intent, without harm, and in a regulatory environment that has been in force for less than two years, faces the same penalty exposure as a repeat violator. The proportionality provisions now in place give competent authorities discretion to calibrate penalties. Discretion is not a right. The EAIF calls on the co-legislators to introduce an explicit remediation window for first-time compliance failures made in good faith: a defined period to correct the violation before penalties apply. This mechanism is standard practice in several Member State administrative traditions. It would reduce the deterrent effect of regulatory uncertainty on innovation without weakening enforcement against deliberate or repeated violations.

Art. 71(6) AI Act – Reduced Fine Ceilings Extended to SMCs

The “whichever is lower” fine ceiling, which uses the lower of a turnover-based percentage or a fixed absolute amount, now applies to SMCs as well as SMEs. The logic is straightforward: a €30 million fine hits a company with €50 million turnover entirely differently from one with €10 billion. That logic does not stop at the SME threshold, and the rules should not either.

Companies that have recently scaled beyond SME status are structurally and financially different from large multinationals even if they no longer qualify as SMEs. Treating them identically for sanctioning purposes creates a penalty cliff that serves no regulatory purpose. The amendment smooths that curve in a targeted way, while preserving the full regime for large enterprises.

For founders and investors in scaling AI companies, this reduces a significant tail risk: the threat of disproportionate fines during periods of rapid growth, when operational and compliance processes are still being built out. Reducing that risk makes the EU a more viable location for AI company formation and scaling without compromising the robustness of enforcement for larger actors.

Art. 2 / Art. 3 AI Act – Intra-Group Deployment: Clarifying the Scope of “Placing on the Market”

The AI Act’s current definition of “placing on the market” does not explicitly exclude intra-group deployments. The result is that a company making an AI system or general-purpose AI model available to other legal entities within the same corporate group may inadvertently trigger full provider obligations, including conformity assessment, registration, and technical documentation requirements designed for external market distribution. This outcome is disproportionate and inconsistent with the underlying purpose of EU product safety law, which is designed to regulate external market access, not internal organisational structures.

Parliamentary amendments⁵ have proposed correcting this directly by clarifying that making an AI system available to other entities within the same corporate group does not constitute “placing on the market.” The EAIF supports this clarification and calls on the co-legislators to adopt it. Internal roll-outs, shared service structures, and group-wide AI deployments are normal features of how modern companies operate. Treating them as market distribution events creates compliance friction with no corresponding safety benefit, and disproportionately affects the increasingly common organisational pattern of vertically integrated AI development across multiple legal entities within a single group. The prohibition framework in Art. 5 must continue to apply regardless of intra-group status; the clarification concerns provider obligations, not the scope of prohibited practices.

Art. 27 AI Act – FRIA and DPIA

Art. 27 requires deployers of high-risk AI systems to conduct a Fundamental Rights Impact Assessment (FRIA). For many deployers, particularly public authorities and organisations already subject to GDPR, this obligation substantially overlaps with the Data Protection Impact Assessment (DPIA) they are already required to carry out under Art. 35 of the GDPR. The current text provides that where FRIA obligations are already met through a DPIA, the FRIA shall “complement” the DPIA, but does not specify how.

Article 27(4) already points in the right direction: where a DPIA has been conducted, the FRIA is required only to complement it, not to replicate it in full. The practical problem is

⁵ (including AM 118, Voss)

that "complement" is undefined. Without clarity on when a DPIA is sufficient to discharge the FRIA obligation in whole or in part, organisations default to conducting two separate assessments covering substantially the same ground, generating administrative cost without generating additional protection.

This problem is part of a broader structural challenge that the Digital Omnibus only partially addresses. For a single AI deployment, an organisation may currently be required to conduct a DPIA, a Transfer Impact Assessment, a Legitimate Interest Assessment, and a FRIA. While the specific legal questions in each framework differ, the foundational factual mapping, covering system architecture, data flows, security protocols, and operational context, remains substantially the same across all of them. Organisations are effectively required to populate the same technical information multiple times across separate regulatory silos, a cumulative burden that falls most heavily on SMEs and startups without dedicated compliance teams.

The EAIF calls on the co-legislators to address this gap by moving towards a "fill in once" approach for risk assessments. A standardised, EU-wide template that allows the core technical information to be recorded once and referenced across multiple regulatory frameworks would substantially reduce duplication without compromising the substance of any individual assessment. The Commission's proposed standardised DPIA template is a welcome step in this direction, but it should be designed with cross-regulatory interoperability in mind from the outset, so that the same foundational information can serve as the basis for FRIA compliance as well.

The EAIF supports the parliamentary amendments proposing that where a DPIA has been conducted and its scope and content are equivalent to what the FRIA requires, the DPIA should be able to fulfil the FRIA obligation in whole or in part. Integration, not parallel duplication, is the correct principle. The EAIF does not support the complete deletion of Art. 27, as proposed by some stakeholders. Fundamental rights assessment serves a distinct and legitimate purpose. What is needed is clarity on how it relates to existing GDPR obligations, so that organisations can meet both requirements through a single, well-designed process rather than two separate administrative exercises covering identical ground.

Annex I Sectors – Conformity Assessment Through Sectoral Authorities

The EAIF supports preserving the current structure of Annex I. AI systems embedded in products covered by Union harmonisation legislation under Section A must remain within the direct scope of the AI Act. A shift to Section B would remove the binding character of AI-specific requirements and replace legal certainty with an open-ended obligation of indeterminate content and timing.

The practical consequences for AI providers are concrete. Under Section A, providers can build on common quality management systems, standardised technical documentation, and horizontal compliance structures that scale across product lines and sectors. Under Section B, they face reduced choice of conformity assessment bodies, longer approval timelines, and sector-specific interpretations of core AI Act concepts such as accuracy, robustness, and transparency. For SMEs and startups operating across multiple regulated product markets, this shift from scalable horizontal compliance to bespoke sector-by-sector approval management is not an abstraction, it is a direct increase in cost, delay, and legal uncertainty at precisely the point where these companies are most vulnerable.

The conformity assessment pathway under Article 43 reinforces this concern. Section A preserves a standards-first compliance route that rewards investment in harmonised standards and allows providers to demonstrate conformity efficiently as standards mature. Section B tends to lock providers into mandatory authority-led review regardless of how well-developed the relevant standards are, making the pathway slower and more expensive without a corresponding gain in safety outcomes for systems that already meet harmonised requirements.

Overlaps and inconsistencies between the AI Act and sectoral legislation should be resolved through timely interpretative guidance: joint implementation notes with sectoral regulators, horizontal guidance documents clarifying the interaction between the AI Act and sectoral frameworks, and practical templates for integrated conformity assessments. For sectors with established specialist regulators (particularly aviation (EASA), automotive, and rail (ERA)) the Commission should clarify how existing sectoral expertise and assessment infrastructure can be used within the AI Act framework, without structural changes to Annex I. The goal is a single, coherent compliance pathway per product, conducted by the right authority, under the full AI Act framework.

Art. 51(2) AI Act – GPAI Systemic Risk Classification: Moving Beyond the Compute Threshold

Art. 51(2) establishes a rebuttable presumption of systemic risk based on a single metric: training compute exceeding 10^{25} floating-point operations. This threshold was already technologically questionable at the time of drafting. It is now demonstrably inadequate. Models trained on high-quality data with efficient architectures achieve capabilities comparable to far larger models trained with substantially more compute. A regulatory trigger calibrated to compute volume does not measure what it purports to measure, and it was designed around a snapshot of model architecture that the industry has already moved beyond. The result is a framework that is obsolete by design: requiring constant revision to remain relevant, and liable to misclassify in both directions, capturing models that pose no systemic risk while missing those that do.

The EAIF calls on the co-legislators to remove Art. 51(2) and make the qualitative and quantitative criteria in Annex XIII the primary basis for systemic risk classification under Art. 51(1)(b). This produces a more technology-neutral framework that weighs capabilities and real-world impact alongside training metrics, rather than relying on a single compute threshold as the sole trigger. To ensure this does not create interpretive uncertainty, the AI Board should be mandated to publish regularly updated guidance specifying indicative thresholds and criteria for each of the Annex XIII factors. The Board is better placed than the legislative text to track technical developments and adjust reference values accordingly. The result would be a framework that remains calibrated to actual risk as the technology evolves, without requiring legislative amendment every time training efficiency improves.

Art. 15, 42(2), 55(1)(d) AI Act – AI Safety and Cybersecurity: Avoiding Regulatory Overlap

The AI Act's provisions on robustness, cybersecurity, and security (Art. 15, 42(2), and 55(1)(d) AI Act) partially overlap with obligations already established under the NIS2 Directive and the Cyber Resilience Act. For companies deploying AI in networked or digital product environments, this creates a compliance landscape in which the same security properties must be assessed, documented, and demonstrated under multiple parallel legal frameworks. The result is increased compliance cost without a corresponding increase in security outcomes.

The EAIF calls on the co-legislators to refine Arts. 15, 42(2), and 55(1)(d) so that they address only AI-specific security requirements, namely those properties that are genuinely distinctive to AI systems and not already covered by NIS2 or the Cyber Resilience Act. General cybersecurity obligations should remain with those instruments, which are designed and resourced to address them. Drawing this boundary clearly reduces duplication, eliminates interpretive ambiguity about which framework governs a given security requirement, and allows companies to manage AI safety and cybersecurity compliance through coherent and non-overlapping processes.

The Parliament's proposed amendment to Art. 15 takes this logic one step further: systems that demonstrably comply with the Cyber Resilience Act should be deemed automatically compliant with Art. 15 of the AI Act. The EAIF strongly supports this approach. Automatic recognition of CRA compliance eliminates the most significant source of duplicative assessment in the cybersecurity overlap and gives product manufacturers a clear, workable pathway. The co-legislators should adopt it.

Art. 111(2) AI Act – Transitional Measures for Legacy AI Systems

AI systems that were lawfully on the market before the AI Act’s application date benefit from transitional protection: they only need to comply with the new framework if they undergo substantial design changes. Requiring retroactive compliance from systems built and assessed under entirely different rules would impose compliance costs with no corresponding safety benefit and would undermine the reasonable expectations of providers who invested in good faith.

The trigger for full compliance (“substantial modification” rather than mere continued operation) is the right line. It protects unmodified legacy systems while ensuring genuinely redesigned ones come under the new framework. The 2030 deadline for public sector systems provides a clear endpoint without disrupting ongoing deployments.

The six-month window for generative AI systems to comply with Art. 50(2) labelling and watermarking requirements are calibrated correctly. Implementing detection and marking infrastructure at scale takes time and a defined grace period prevents market disruption while maintaining pressure to act. At six months, this is an implementation accommodation, not a substantive exemption. One structural improvement the EAIF calls for: the transitional protection afforded to legacy systems should be embedded directly in the operative provisions of Art. 111, not left to rely on interpretive inference from transitional clauses. Legal certainty requires that providers can identify their transitional status from the face of the operative text without cross-referencing recitals or implementation guidance. Where the protection is substantive, the operative article should say so explicitly.

Art. 113(3)(d) AI Act – Conditional Commencement of High-Risk AI Obligations

The most significant practical risk in the AI Act’s implementation timeline is that binding obligations take effect before the tools companies need to comply with them actually exist. Harmonised standards were still being developed when the Act was finalised. Guidance, supervisory capacity, and conformity assessment infrastructure are still maturing. Art. 113(3)(d) addresses this directly: the most demanding Chapter III obligations only start once the Commission has formally confirmed that essential prerequisites are in place.

From that Commission decision, staggered transition periods apply: six months for Annex III stand-alone systems, twelve months for Annex I product-embedded systems. The differentiation reflects real differences in integration complexity and assessment pathways.

The EAIF calls on the co-legislators to write fixed dates directly into the operative text: 2 December 2027 for Annex III systems and 2 August 2028 for Annex I. These dates must be treated as firm outer limits, not as aspirational backstops. The Commission must be formally obliged to publish the necessary harmonised standards and guidance before those dates. Where that obligation cannot be met, the dates should shift accordingly, not through discretionary Commission action, but through an automatic and transparent adjustment mechanism tied to the publication of standards. A fixed, publicly known commencement date, backed by a delivery obligation on standards, is the single most effective tool for enabling compliance planning: without it, companies cannot make informed investment decisions about staffing, tooling, or conformity assessment. The argument for realistic deadlines is not about lowering standards; it is about ensuring that when obligations apply, companies are actually in a position to meet them, because standards exist, supervisory structures are functional, and conformity assessment bodies are available. The Commission-trigger mechanism as currently drafted, under which obligations begin to run only after the Commission determines that appropriate measures in support of compliance are in place, introduces a new source of legal uncertainty: companies cannot plan around a decision whose timing is entirely within the Commission's discretion. Fixed dates with a corresponding standards delivery obligation are the only design that provides actionable planning certainty.

Further Welcomed Amendments

Several further amendments in the package address technical and procedural gaps that the EAIF welcomes without requiring detailed comment. The consolidation of bias-testing data rules into Art. 4a, with the corresponding removal of the overlapping Art. 10(5), improves the coherence and accessibility of the framework without altering its substance. The explicit permission to reuse existing conformity assessment documentation for AI Act notification purposes under Art. 29(4) removes a procedural redundancy that served no oversight purpose. The standardisation of notification categories via the Annex XIV coding system under Art. 30(2) eliminates a source of cross-border inconsistency in notified body designation. And the introduction of a flat-rate fee structure for Member State consultations with the Scientific Panel under Art. 69(2) and (3) removes cost as a factor in whether national authorities seek expert input, which benefits the quality and consistency of enforcement across the single market. These are proportionate, well-targeted measures and the EAIF supports their adoption.

Broader Context of the Digital Omnibus

The EAIF broadly welcomes the Digital Omnibus on AI. Taken together, the amendments represent a coherent effort to make the AI Act's implementation framework more proportionate, more certain, and more operationally usable, particularly for European

SMEs, startups, and SMCs. The Commission's projected savings of up to €1 billion annually, cumulating to approximately €5 billion by 2029⁶, reflect the genuine scale of the compliance burden these reforms address. For European AI companies competing globally, where regulatory overhead directly affects investment and speed-to-market, these gains are real.

Reducing unnecessary administrative overhead for innovative companies is sound policy. Diluting substantive rights or protections without clear justification is not, and the European Parliament should insist on that distinction. Where amendments shift the balance of obligations or enforcement safeguards, the case for change must be demonstrated. Institutional capacity must also keep pace: the AI Office and national competent authorities need adequate resources to discharge their expanded responsibilities, since simplification that outpaces enforcement capacity trades one form of uncertainty for another. One specific structural gap requires attention in this regard: the AI Office should be given a formal mediation and harmonisation role for cases where national competent authorities arrive at divergent interpretations of the same provision. As enforcement matures across 27 Member States, interpretive fragmentation is a predictable and serious risk, and companies operating across borders cannot manage compliance effectively when the same system is assessed differently in different jurisdictions. A mandatory consultation mechanism, with a defined timeline and a requirement that the AI Office issue a binding or advisory opinion, would substantially reduce that risk.

The Digital Omnibus will ultimately be judged by what it delivers in practice: a framework that European AI companies can actually navigate, obligations that take effect on fixed, publicly known dates backed by a corresponding delivery obligation on standards, and an enforcement architecture calibrated to the reality of European AI development rather than to conditions that do not yet exist. Achieving that requires the co-legislators to move with purpose and without adding further layers of complexity to a framework that is already demanding. The goal is a regulatory environment in which European AI companies can build, scale, and compete, and in which the protections the AI Act was designed to deliver are real rather than nominal.

⁶ https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2718.

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 3000 members, making us the largest European AI organisation. We represent members such as SMEs, corporates, organisations and individual experts in AI.

Contact:

info@eaiforum.org