

Position Paper on the European Commission's Public Consultations concerning the Digital Omnibus Package

Opening Remarks

Employers of Poland welcomes the European Commission's Digital Omnibus initiative as a timely and necessary effort to streamline the EU's digital regulatory framework. This position is informed by the substantive input and practical insights of our member companies, who are Polish leaders in logistics, e-commerce, and digital services operating across multiple Member States. By deploying AI-driven solutions at scale, specifically in areas such as route optimization, parcel management, and delivery forecasting, as well as customer service automation within both the logistics and retail and consumer goods sectors - these enterprises provide our organization with a direct and practical stake in the successful implementation of the AI Act.

We recognize and appreciate the Commission's intent to reduce compliance burdens, improve legal certainty, and align the AI Act's implementation timeline with the actual readiness of the standards and guidance ecosystem. Furthermore, the proposed grace periods for high-risk systems, the shift towards encouraged rather than prescriptive AI literacy, and the strengthening of the AI Office's central oversight role are all considered steps in the right direction.

However, it is our assessment that the current proposals do not go far enough. Several critical gaps remain that, if left unaddressed, risk perpetuating legal uncertainty and fragmenting enforcement, which would ultimately undermine the very competitiveness and innovation objectives that the Omnibus is designed to support.

Employers of Poland remains committed to responsible AI development and deployment, as well as to constructive engagement with EU institutions as the Digital Omnibus progresses through the legislative process. The recommendations set out below are offered in that spirit, with the primary aim of ensuring that the AI Act delivers on its promise of trustworthy AI while enabling European businesses to innovate, compete, and grow in a global market.

1. Coherent Timelines for the AI Act: Grace Periods, Fixed Deadlines, and Transparency Rules

The Digital Omnibus rightly acknowledges that the AI Act's implementation timeline must be recalibrated. However, the current proposals address this need in a fragmented manner, extending relief to some obligations while leaving others enforceable under conditions of equivalent uncertainty. For our member companies, who deploy AI-driven solutions across logistics, e-commerce, and digital services spanning multiple Member States, this piecemeal approach complicates compliance planning and undermines confidence in the regulatory process. Three interconnected gaps must be addressed in the final legislative text.

First, the grace period must extend beyond the High-Risk regime to encompass the GPAIM provisions. Despite GPAIM obligations being enforceable since 2 August 2025, the foundational conditions for compliance, finalised standards, clear guidance, and operational enforcement infrastructure, remain absent. These are precisely the deficiencies that justified deferring the High-Risk regime. Maintaining a grace period for one set of obligations while enforcing another subject to identical implementation gaps creates regulatory incoherence. For our members, who integrate general-purpose AI capabilities into logistics operations and customer-facing services as deployers rather than model developers, the GPAIM framework's blurring of the boundary between model development and system deployment is a direct source of legal uncertainty that the Omnibus must resolve.

Second, conditional and staggered timelines for high-risk systems should be replaced with a single, fixed deadline. The current mechanism, contingent on the Commission's confirmation of ecosystem readiness, with fallback deadlines of December 2027 for Annex III and August 2028 for Annex I - generates uncertainty that enterprises cannot plan around. The criteria triggering confirmation remain ambiguous, and the staggered dates add complexity without policy justification. A single, unified deadline for all high-risk systems under both Annexes would deliver the legal certainty that businesses need for effective compliance planning. That deadline must be realistic: with harmonised standards not expected before the end of 2026, the lead time must allow for genuine implementation rather than paper compliance. Employers of Poland does not advocate for indefinite delay, we advocate for a single, credible deadline that businesses can plan around with confidence.

Third, the transparency grace period under Article 50 must cover all provisions and all AI systems in scope. The current proposal grants providers of generative AI systems just six months to comply with Article 50(2), and only for systems on the market before 2 August 2026. This narrow scope is unjustified. The technical challenges of implementing marking and labelling requirements are universal, not system-specific, and the standards ecosystem remains under active development, the second draft Code of Practice on Marking and Labelling was published as recently as March 2026. Systems placed on the market after the cut-off date

would be expected to comply with requirements derived from a Code finalised only weeks earlier, an expectation that is neither realistic nor conducive to robust implementation. The Commission must extend the grace period to all provisions of Article 50, with a single, unified compliance deadline of February 2027 for all AI systems in scope, irrespective of market entry date.

Taken together, these three measures would ensure that the AI Act's transition is coherent, predictable, and grounded in the actual readiness of the regulatory ecosystem, rather than imposing obligations that businesses cannot yet meaningfully fulfil.

2. Explicit Pro-Innovation Mandate for Regulators

While the Omnibus proposal represents a step towards a more centralised approach to the AI Act, it does not provide a clear mandate for regulators, at either the EU or national level, to actively champion innovation. This gap means that no digital regulator, including national competent authorities, is explicitly charged with fostering an innovation-friendly environment. Although Article 1 underscores the importance of human-centric and trustworthy AI, this objective is insufficiently reflected in the mandates or competencies assigned to regulatory bodies. Without an explicit innovation mandate, the Act risks giving rise to inconsistent interpretation and regulatory fragmentation, which would undermine effective implementation and weaken the EU's global competitiveness in AI. To drive investment, innovation, and sustainable growth, digital regulators must share a unified mission to cultivate a stable and forward-looking regulatory environment.

Regulators responsible for the implementation and enforcement of the AI Act should be provided with a clear and explicit mandate to support innovation alongside the protection of fundamental rights. This mandate should:

- Affirm that innovation, competitiveness and fundamental rights protection are complementary objectives.
- Encourage proportionate and risk-based enforcement.
- Promote legal certainty and predictability in supervisory practice.
- Require regulators to consider economic impact and technological feasibility when interpreting ambiguous provisions.

A pro-innovation mandate should not weaken safeguards. Rather, it should ensure that enforcement focuses on material risks rather than formalistic compliance, supervisory expectations are aligned with technological realities, and novel AI applications are not discouraged due to regulatory uncertainty. Such a mandate would also help avoid divergent enforcement cultures across Member States, thereby strengthening the integrity of the internal market.

Commission's proposal does not yet address this gap. Amendments to the text remain largely confined to enforcement powers and procedural timelines, without equipping regulators with an explicit mandate to promote innovation. This omission is inconsistent with the broader stated objectives of advancing EU competitiveness and digital sovereignty. The final legislative text should expressly task national competent authorities and the AI Office with fostering innovation alongside the protection of fundamental rights, ensuring that regulatory practice reflects the full scope of the AI Act's objectives.

3. Enabling Effective Bias Detection Across All AI Systems

The Omnibus proposal's introduction of Article 4a marks a significant and necessary shift, broadening the legal basis for processing Special Categories of Data (SCD) to detect and correct bias across all AI systems and models, not just those deemed high-risk. This is a crucial improvement that will empower more effective AI training and support the EU's commitment to fairness and responsible innovation. It acknowledges a fundamental reality: fairness cannot be meaningfully tested without appropriate reference data, and in many contexts bias manifests precisely along protected characteristics. Extending the legal pathway to process SCD, under strict safeguards, supports more robust model evaluation, improves accountability, and strengthens the EU's commitment to responsible innovation.

For our member companies operating in logistics, e-commerce, and digital services, the ability to identify and address discriminatory patterns is not a peripheral compliance exercise, it is integral to the responsible deployment of AI systems that affect customers, employees, and service partners on a daily basis. Restricting bias mitigation to high-risk systems would leave a substantial portion of the AI ecosystem untested for discrimination, creating a regulatory blind spot that is inconsistent with the AI Act's stated objective of preventing harm. Bias is a structural feature of data-driven systems, not a function of their risk classification. Every model must be capable of detecting, measuring, and correcting discriminatory outcomes, and this requires access to SCD under appropriately strict safeguards.

If the final Omnibus text does not retain this approach, organisations will be pushed into one of two undesirable outcomes: (i) under-testing models for bias (increasing discrimination risk), or (ii) using uncontrolled proxies (increasing both discrimination and data protection risk). The proposed approach should remain in the final text to enable Europe's AI ecosystem, including the TSL and e-commerce sectors, to innovate responsibly while maintaining strong safeguards for individuals.

The possible introduction of a "strict necessity" standard for processing special categories of personal data for bias detection, while intended to maintain rigorous standards, risks undermining the provision's intended purpose if applied too restrictively. This elevated threshold, if interpreted too narrowly, could effectively preclude organisations from accessing

the data required to detect, measure, and correct discriminatory patterns, particularly in sectors where bias manifests in subtle and complex ways. To fulfil the objective of responsible AI development, the "strict necessity" standard should be interpreted in a manner that enables effective bias detection rather than constituting a de facto barrier to responsible data use.

4. Undermining the Risk-Based Principle Through Asymmetric Application

The current Omnibus proposal includes an amendment to Article 99, extending existing regulatory privileges on penalties for SMEs to SMCs. This asymmetric introduction of obligations under the AI Act undermines the legislation's fundamental risk-based approach, which is designed to scale requirements according to the level of risk posed by an AI system, not the size or identity of its provider.

Risk is determined by the nature and deployment of the AI system, not by the developer's size. High-risk AI systems can be created by startups as well as large corporations, and the potential for harm is inherent in the system itself. Using company size as a proxy for risk misrepresents the Act's intent and creates regulatory blind spots where robust safeguards are most needed.

To maintain coherence and effectiveness, the AI Act's obligations must apply symmetrically to all entities developing or deploying systems of comparable risk. Deviating from this principle dilutes protections for citizens, creates an uneven market, and ultimately weakens both safety and competitiveness within the Union.

Current legislative developments reflect a more measured approach to regulatory simplification by maintaining certain obligations, such as registration requirements for specific AI systems, which reinforces the principle that risk-based obligations should not be attenuated based on provider characteristics. However, the penalty asymmetry for SMCs remains a significant concern, as it continues to distort the risk-based framework. The final text should ensure that all entities developing or deploying systems of comparable risk are subject to equivalent obligations and proportionate enforcement.

5. Equipping the AI Office for Effective Cross-Border Enforcement (Article 75)

Employers of Poland supports the consolidation of supervisory authority within the AI Office as a necessary condition for regulatory coherence across the Union. For our member companies, operating logistics networks, e-commerce platforms, and digital services that span multiple Member States, the prospect of navigating divergent national enforcement regimes represents a material compliance risk and a significant barrier to cross-border scaling.

The proposed amendments to Article 75, which designate the AI Office as the central enforcement authority for AI systems integrated within Very Large Online Platforms (VLOPs) and Very Large Online Search Engines (VLOSEs), as well as for GPAIMs, represent a sound structural choice. Aligning supervisory architecture with the model already established under

the Digital Services Act reduces the potential for jurisdictional overlap and ensures that platforms with Union-wide reach are subject to consistent standards rather than a patchwork of national interpretations.

However, centralisation must be accompanied by operational credibility. The AI Office's expanded remit will require it to develop deep sectoral expertise, including in areas such as logistics, transport, and e-commerce, where AI systems are embedded in time-sensitive, high-volume operations. To this end, we recommend that the final legislative text:

- Establish clear mechanisms for the AI Office to monitor specific AI systems on an ongoing basis.
- Ensure that the AI Office is resourced commensurately with its expanded mandate, including through the recruitment of technical specialists with relevant industry experience.
- Create structured channels for industry consultation, enabling the AI Office to ground its supervisory practice in operational realities and to anticipate implementation challenges before they become enforcement disputes.

Without these safeguards, the centralisation of oversight risks becoming an institutional aspiration rather than a practical reality. The final text must ensure that the AI Office is not merely designated as the central authority, but equipped to function as one.

6. Context-Sensitive AI Literacy: From Mandates to Meaningful Competence

The shift from prescriptive AI literacy obligations to an encouragement-based model is a pragmatic recognition of the diversity that characterises Europe's business landscape. Employers of Poland welcomes this recalibration, which reflects legitimate concerns raised by our member companies regarding the impracticability of uniform literacy mandates across organisations that differ vastly in size, sector, and technological maturity.

In our members' experience, AI literacy is most effective when it is tailored to operational context. A warehouse operative interacting with an AI-driven parcel sorting system requires a fundamentally different level and type of understanding than a data scientist developing the underlying model or a compliance officer assessing its risk profile. Prescriptive, one-size-fits-all requirements risk producing compliance exercises that satisfy regulatory form without delivering genuine comprehension.

That said, encouragement must not become an invitation to inaction. We recommend that the final legislative text be supplemented with practical Commission guidance on what meaningful AI literacy looks like across different organisational functions and risk levels, drawing on sectoral best practices. In particular, guidance should address:



- How AI literacy expectations scale with the risk classification of the systems in question, ensuring that operators of high-risk systems invest proportionately in staff competence.
- The distinction between general AI awareness (appropriate for all staff interacting with AI outputs) and technical AI competence (necessary for those designing, deploying, or overseeing AI systems).
- The role of ongoing learning, recognising that AI capabilities and risks evolve rapidly and that a one-time training exercise is unlikely to remain adequate over a system's lifecycle.

This approach preserves the flexibility that businesses need while establishing a clear expectation that AI literacy is a shared responsibility, one that Member States, the Commission, and industry must pursue collaboratively, with the seriousness that the deployment of consequential AI systems demands.

7. Watermarking Obligations (Article 50)

7.1. Allocating Watermarking Obligations to Platform and Framework Providers

The transparency obligations under Article 50, particularly those requiring AI system outputs to be marked in a machine-readable format and detectable as artificially generated or manipulated, raise a critical question of regulatory design regarding where in the value chain should the obligation to implement watermarking mechanisms be placed.

For our member companies, many of whom deploy AI agents built on third-party agentic development kits (ADKs) and multi-agent orchestration platforms, this question has immediate practical significance. A business user configuring an AI agent to generate logistics reports, customer communications, or claims summaries is typically not in a position, technically or organisationally, to implement output watermarking at the system level. Requiring such users to do so would be disproportionate and, in many cases, technically unworkable.

The obligation to embed watermarking mechanisms must therefore fall on the providers of the underlying platforms and frameworks that facilitate the creation of AI agents. Where an ADK or agentic AI framework enables the generation of content, it is the framework provider that controls the technical infrastructure necessary to implement robust, consistent watermarking across all outputs. Placing this obligation on individual agent creators, who may lack programming expertise entirely, would fragment implementation, reduce watermark reliability, and impose compliance costs that bear no reasonable relationship to the user's role in the value chain.

7.2. Applying Proportionality to Text Watermarking

Watermarking obligations for text-based AI outputs must be subject to a meaningful proportionality assessment. Not every text generated by an AI system carries the same risk of misleading its audience, and the regulatory framework should reflect this reality.

In the operations of our member companies, AI agents routinely produce internal-facing content. These outputs are consumed within controlled environments where the provenance of the content is known and documented through the organisation's IT and data architecture. Mandating machine-readable watermarks for such outputs would impose technical costs without corresponding benefit, particularly where the organisation's technical and data architecture already ensures traceability of AI-generated resources through system-level metadata, access controls, and audit trails or the outputs are intended exclusively for internal use and are accompanied by clear, human-readable indications that they were generated with the assistance of AI.

The case for watermarking is materially different when AI-generated text enters the public domain. In these contexts, the risk of audiences being misled about the origin of the content is genuine, and technical watermarking serves a clear and proportionate purpose.

We therefore recommend that the final legislative text distinguish explicitly between these use cases, ensuring that watermarking obligations are calibrated to the actual risk of deception rather than applied indiscriminately to all AI-generated text. This distinction is particularly important for sectors such as logistics and e-commerce, where AI-generated internal documentation is a routine operational tool and blanket watermarking requirements would impose a disproportionate administrative burden with no meaningful gain in transparency for end users or the public.

8. Reconciling Disclosure with Trade Secret Protection

The European Commission should strive to ensure that documentation obligations do not force companies to disclose key algorithms that constitute their know-how. Transparency should not be interpreted as an obligation to reveal "how the model works" in a way that enables reverse engineering. Instead, EU guidance and harmonised standards should promote a layered transparency approach: meaningful explanations for regulators and affected individuals, while preserving trade secrets and security-sensitive details.

Efforts should be made to develop standards that reconcile transparency with the protection of intellectual property and confidential business information, including:

- **Role-based disclosure:** different granularity for (i) competent authorities/auditors, (ii) business customers/partners, and (iii) end users/consumers.



- **Evidence-based transparency:** focus on outcomes and controls (risk management, performance metrics, known limitations, bias testing results, monitoring, incident handling) rather than disclosure of source code or detailed model weights.
- **Secure audit mechanisms:** controlled access for regulators (e.g., onsite inspection, secure data rooms, confidentiality undertakings), allowing verification without publication.

To reconcile transparency and intellectual property protection, we propose a structured disclosure model:

- **Level 1 – User Transparency:** Clear, understandable information on system purpose, logic, and impact.
- **Level 2 – Regulatory Access:** Full technical documentation available to competent authorities.
- **Level 3 – Controlled Audit:** Secure, supervised audit environments for independent review where required.
- **Level 4 – IP Protection:** No obligation to disclose proprietary weights, source code, or trade secrets beyond what is strictly necessary for the purposes of fairness and transparency of AI systems.

This approach ensures compliance while preserving market competitiveness of entities engaged in creating innovative approaches to their business and social activity.

From a personal data protection perspective, this balance is particularly important. Data subjects must receive understandable information about the logic involved and the significance/consequences where relevant, but GDPR does not require disclosure of trade secrets or full algorithmic detail. In practice, companies should be able to comply through high-level, intelligible explanations (what factors matter, what safeguards exist, how to contest outcomes, how human oversight works), without revealing proprietary methods.

This is particularly relevant in logistics and e-commerce, where AI-driven components (ETA prediction, route optimisation, fraud detection, parcel prioritisation, automated customer support and claims triage) could be exploited if overly detailed logic were disclosed, creating security risks (gaming fraud controls, manipulating delivery prioritisation) and undermining service integrity. Standards should therefore explicitly recognise the need to protect security-by-design and anti-abuse measures alongside trade secrets.

Finally, the Omnibus "consumer fairness" perspective reinforces that transparency should be clear and non-misleading, not necessarily "fully technical". The objective should be to prevent deceptive practices and enable informed choices and effective redress (complaints, human

review, appeal), while ensuring that intellectual property and confidential operational methods remain protected.

9. Practical Definitions for a Workable AI Act

The AI Act's implementation is currently undermined by a significant gap between the legislation's ambitions and the clarity of its provisions. For our member companies, including large, multinational enterprises coordinating AI-driven operations across multiple Member States, this ambiguity is not an abstract concern. It directly impedes internal compliance planning, forces costly conservative assumptions, and in some cases risks paralysing the very deployment of AI systems that the Act seeks to regulate responsibly. Three categories of difficulty stand out.

First, key definitions remain too vague to support consistent application. The concept of "AI system" itself, the threshold of "significant generality" for general-purpose AI models, and related terms lack the precision needed for businesses to determine with confidence whether and how a given system falls within scope. For our members, the boundary between a conventional software system and an AI system subject to the Act's obligations is frequently unclear. This uncertainty is compounded by the breadth of the definitions, which risk capturing well-established algorithmic tools that were never contemplated as targets of AI-specific regulation.

Second, certain technical requirements are unrealistic in practice. The obligation to measure FLOPS spent on model training, particularly for fine-tuned models built on third-party foundations, cannot be reliably implemented across the industry. Our member companies, as deployers rather than foundational model developers, often lack visibility into the computational resources consumed during upstream training. Enforcement of such provisions risks producing either unreliable self-reporting or strategic avoidance, neither of which serves the Act's objectives.

Third, internal inconsistencies between the Act's legislative purpose and its operative provisions, as well as between the Act and related official guidance, create confusion that regulatory bodies have not yet resolved. The suggestion that a single-modality model may qualify as a general-purpose model is one such example. For companies attempting to build compliant governance frameworks, contradictory signals from legislative text, recitals, and Commission guidance make it impossible to identify a single authoritative interpretation.

To address these issues, the final legislative text and accompanying guidance must be aligned into a single, coherent framework. All definitions and provisions should be reviewed for practical applicability, with consideration given to removing criteria that are difficult to measure or enforce - particularly those related to systemic risk thresholds. Most importantly, a clear and binding implementation roadmap must be published and adhered to by all

regulatory bodies, providing the stability that businesses require to invest in compliance with confidence rather than hedging against regulatory ambiguity.

10. Reducing Classification Uncertainty for AI Systems

The risk-based classification of AI systems, into unacceptable, high, limited, and minimal risk categories, is a sound regulatory principle. In practice, however, the boundaries between these categories remain insufficiently defined, creating one of the most significant compliance challenges facing our member companies. It is crucial that the Commission publishes detailed interpretative guidelines, supplemented by sector-specific case studies, that clarify the practical application of classification criteria to real-world AI systems. For detailed recommendations on sector-specific guidance and case studies, see Section 11.5.

This need is even stronger when taking into account the combined compliance burden: AI Act duties (risk management, data governance, technical documentation, human oversight) must be aligned with GDPR principles (lawfulness, purpose limitation, data minimisation, transparency, storage limitation, accountability), and with Omnibus-driven expectations on fair consumer information and avoiding misleading practices or "dark patterns" in digital journeys.

From a personal data perspective, classification uncertainty directly impacts:

- Whether the company must perform DPIAs (and in some cases consult the authority) due to high risks to individuals.
- How to implement privacy by design/by default (e.g., minimising tracking and profiling in operations).
- Whether certain uses may trigger GDPR Article 22 concerns (automated decision-making with legal or similarly significant effects), particularly when AI impacts customers' access to services (e.g., fraud flags leading to blocking), employees' task allocation, or couriers' performance scoring.

For the logistics, TSL, and e-commerce sectors in particular, standard optimisation systems, such as route planning, demand forecasting, and warehouse allocation, should be excluded from the high-risk category where they do not directly determine decisions on hiring, dismissal, or pay conditions, and where the final decision authority rests with a human operator such as a dispatcher or operations manager.

Furthermore, AI systems are subject to continuous updates and retraining cycles, yet the current framework leaves businesses uncertain about which modifications trigger a full conformity assessment for high-risk systems. This ambiguity discourages iterative improvement, as routine enhancements to accuracy or safety could expose companies to disproportionate reassessment costs. This uncertainty also affects data protection

governance: updates may change the purposes and means of processing, alter the risk profile, or expand categories of personal data used (e.g., adding geolocation granularity, behavioural signals, courier telemetry), potentially requiring DPIA updates and revisiting transparency notices and retention schedules.

For detailed recommendations on a materiality-based approach to updates, see Section 11.4.

11. Strengthening Fairness, Non-Discrimination and Responsible AI Design vs. Ensuring Innovation-Supportive AI Governance

To ensure that regulatory flexibility does not undermine fundamental rights protection, AI systems should be designed, deployed and monitored in a manner that actively safeguards fairness, non-discrimination and transparency. In this context, we recommend the following principles:

11.1. Periodic and Risk-Proportionate Bias Testing

AI systems, particularly those that may affect individuals' rights, economic opportunities, or access to services, should undergo periodic bias testing where appropriate and proportionate to their risk profile. Bias testing should:

- Be tailored to the specific context and function of the system.
- Consider reasonably foreseeable use cases.
- Be conducted at appropriate stages of the lifecycle, including post-deployment where systems evolve.
- Include both quantitative and qualitative assessments where relevant.

The frequency and depth of testing should be risk-based, taking into account the sensitivity of the application domain (e.g. employment, credit, insurance, public services), the scale of deployment and the potential impact on vulnerable groups. This approach supports accountability while avoiding unnecessary burdens for low-risk applications.

11.2. Sector-Specific Fairness Metrics

Fairness is context-dependent. Therefore, AI systems should incorporate fairness metrics suited to their sector and intended purpose. We recommend:

- Development of sector-specific fairness benchmarks (e.g. in logistics, financial services, e-commerce, recruitment, mobility).
- Use of explainable and auditable fairness indicators appropriate to the operational environment.
- Avoidance of a one-size-fits-all metric that may not reflect sector realities.

Fairness metrics should be technically feasible, aligned with anti-discrimination law, interpretable by regulators and auditors, and adaptable to evolving technological standards. Clear guidance on acceptable methodologies would significantly reduce uncertainty and support harmonised implementation across EU Member States.

11.3. Safeguards in Case of Automated Decision-Making

Where AI systems may fall within the scope of Article 22 GDPR (automated decision-making producing legal or similarly significant effects), appropriate safeguards must be embedded by design. Such safeguards should include meaningful human oversight mechanisms, the possibility for individuals to obtain clarification and contest decisions, clear allocation of responsibility within the organisation, and transparent communication about the involvement of automated processing where relevant.

At the same time, human oversight should be effective and not merely symbolic, proportionate to the impact of the decision, and supported by adequate training and authority to intervene. This ensures alignment between AI Act obligations and existing data protection law, while providing legal certainty for businesses.

11.4. Continuous Monitoring and Lifecycle Responsibility

AI governance should reflect the dynamic nature of AI systems. Where systems are updated, retrained, or fine-tuned, organisations should assess whether changes materially affect fairness, risk profile or data protection implications. A materiality-based approach should apply, avoiding unnecessary full reassessment for minor optimisations while ensuring re-evaluation where risk materially increases.

In logistics and e-commerce, a practical and proportionate approach is needed: update-related obligations should be tied to material changes (e.g., new intended purpose, new affected population group, new decision impacts, or significant performance drift) rather than routine retraining. Otherwise, companies may be discouraged from improving accuracy and safety. Sectoral guidelines should therefore include: (i) a materiality test for updates, (ii) examples for ETA/routing/fraud/claims automation, and (iii) minimum expectations for documentation, monitoring, incident handling, and user/customer communications—consistent with the Omnibus goal of preventing misleading practices and ensuring clear information.

11.5. Sector-Specific Guidance and Case Studies

To reduce classification uncertainty and compliance fragmentation across EU Member States, we strongly recommend the development of:

- Sector-specific guidance clarifying high-risk vs. non-high-risk classification.
- Practical case studies illustrating compliant implementation models.

- Interpretative examples regarding significant modification and model updates.
- Harmonised guidance on interaction between AI Act obligations and GDPR requirements.

Such guidance is particularly urgent for sectors including TSL and e-commerce, where detailed interpretations from the European Commission on which AI systems are subject to specific obligations would reduce legal uncertainty and make it easier for companies to prepare for the implementation of the regulations. This need is reinforced by the combined compliance burden: AI Act duties (risk management, data governance, technical documentation, human oversight) must be aligned with GDPR principles (lawfulness, purpose limitation, data minimisation, transparency, storage limitation, accountability), and with Omnibus-driven expectations on fair consumer information and avoiding misleading practices or "dark patterns" in digital journeys.

Such guidance should be developed in cooperation with industry and widely-considered society, publicly accessible and regularly updated, and consistent across EU institutions and national authorities. Clear and practical examples are essential to avoid divergent interpretations and to enable SMEs and SMCs to comply efficiently without disproportionate administrative burden.

11.6. Expansion and Coordination of EU-Level Regulatory Sandboxes

We support the development of coordinated EU-level AI regulatory sandboxes and recommend expanding their scope and accessibility. However, to maximise their effectiveness, sandboxes should:

- Operate under transparent and publicly available selection criteria.
- Provide equal access opportunities to SMEs, SMCs and start-ups.
- Ensure sectoral diversity in participation.
- Include structured cooperation between national authorities and the AI Office.
- Allow cross-border testing where relevant.

Selection criteria should be risk-based and innovation-focused, non-discriminatory and transparent, and clearly linked to public-interest objectives. In addition, sandbox participation should provide regulatory guidance during development and testing, legal clarity on compliance expectations, feedback loops between innovators and regulators, and documented best practices that can inform future guidance.

A well-designed EU sandbox ecosystem can function as a bridge between regulatory objectives and technological experimentation, enabling safe innovation without regulatory arbitrage.

11.7. Harmonised Interpretative Roadmaps

To reduce fragmentation and classification uncertainty, we recommend publishing harmonised and binding interpretative roadmaps at EU level. These roadmaps should clarify timelines for implementation across all major AI Act obligations and provide consolidated interpretations of key concepts (e.g. high-risk classification, significant modification, GPAI scope). They may also prove effective in explaining interactions between the AI Act and other relevant legislation, including GDPR, consumer protection law and sectoral regulations. It is highly advisable that proposed roadmaps also include practical compliance pathways adapted to different categories of operators.

Interpretative roadmaps should be adopted in coordination with the European AI Board and be updated regularly in response to technological developments. That would provide consistent guidance across EU Member States and reduce reliance on divergent national interpretations.

Where appropriate, interpretative documents should include decision trees for classification, illustrative case studies, examples of compliant documentation practices and clear explanations of supervisory expectations.

Binding and harmonised guidance is essential to avoid compliance fragmentation, legal uncertainty and competitive distortions within the EU internal market.

Closing Remarks

The Digital Omnibus represents a welcome acknowledgement by the European Commission that the AI Act's implementation must be recalibrated to reflect the realities facing businesses across the EU. Recent refinements to the legislative proposal mark significant progress, particularly in establishing fixed deadlines for high-risk systems and strengthening the role of centralised oversight. However, as set out in this position paper, significant gaps remain.

Employers of Poland calls on the European Parliament and the Council to address the following key concerns in the final legislative text:

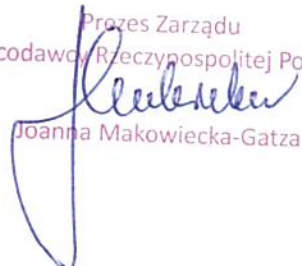
- **The grace period must be extended beyond high-risk systems to include GPAIM provisions**, ensuring regulatory consistency and acknowledging that implementation challenges are not confined to a single risk category.
- **Conditional and staggered timelines for high-risk systems should be replaced with a single, fixed deadline**, providing all stakeholders with the clarity needed for effective compliance planning.



- **The transparency grace period under Article 50 must cover all provisions and all AI systems in scope**, rather than applying narrowly to a single sub-article and pre-existing systems only.
- **Regulators at both EU and national level need an explicit innovation mandate**, ensuring that enforcement is proportionate, forward-looking, and aligned with Europe's competitiveness ambitions.
- **The "strict necessity" standard for processing Special Categories of Data must be interpreted purposively**, enabling genuine bias detection rather than creating barriers to responsible AI development.
- **Classification uncertainty, trade secret protection, and sector-specific guidance** require urgent attention, particularly for sectors such as logistics, TSL, and e-commerce, where AI is integral to daily operations but risk profiles differ significantly from the high-risk use cases the Act primarily targets.

These are not calls for deregulation. They are calls for smarter, more coherent regulation, regulation that is proportionate, predictable, and aligned with the EU's stated objectives of fostering trustworthy AI while safeguarding European competitiveness.

As an organization representing enterprises at the forefront of AI-driven logistics and e-commerce, Employers of Poland stands ready to contribute to this vital effort. We welcome continued dialogue with the European Commission, the European Parliament, the Council, and national competent authorities as the Digital Omnibus advances through the legislative process. It is our firm belief that a well-calibrated AI Act, one that provides legal certainty, supports responsible innovation, and protects fundamental rights is in the interest of all stakeholders, from business operators to end-users. Therefore, we urge the co-legislators to seize this opportunity to get it right, ensuring a regulatory framework that fosters European competitiveness while maintaining the highest standards of safety and trust.

Przes Zarządu
Pracodawcy Rzeczypospolitej Polskiej

Joanna Makowiecka-Gatza