CMS
law·tax·future

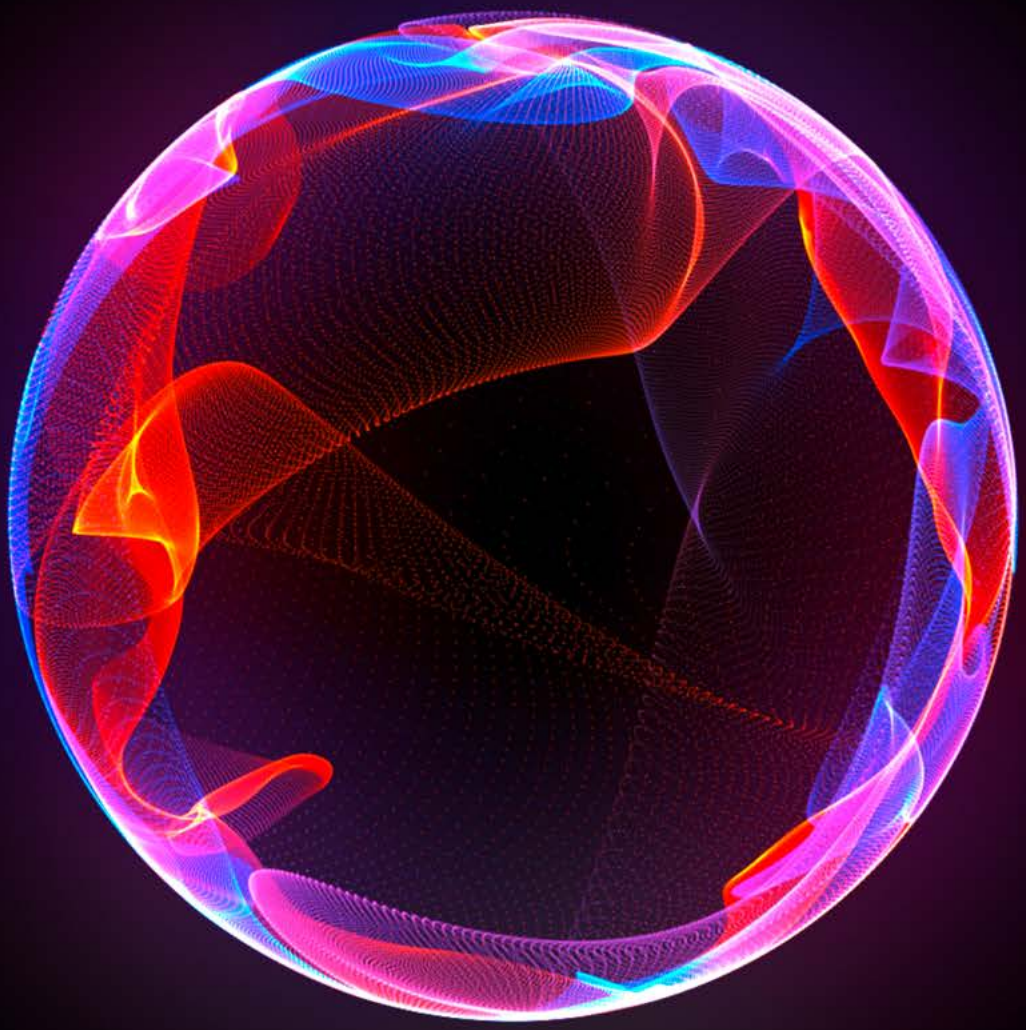

Guide on AI Omnibus

July 2026

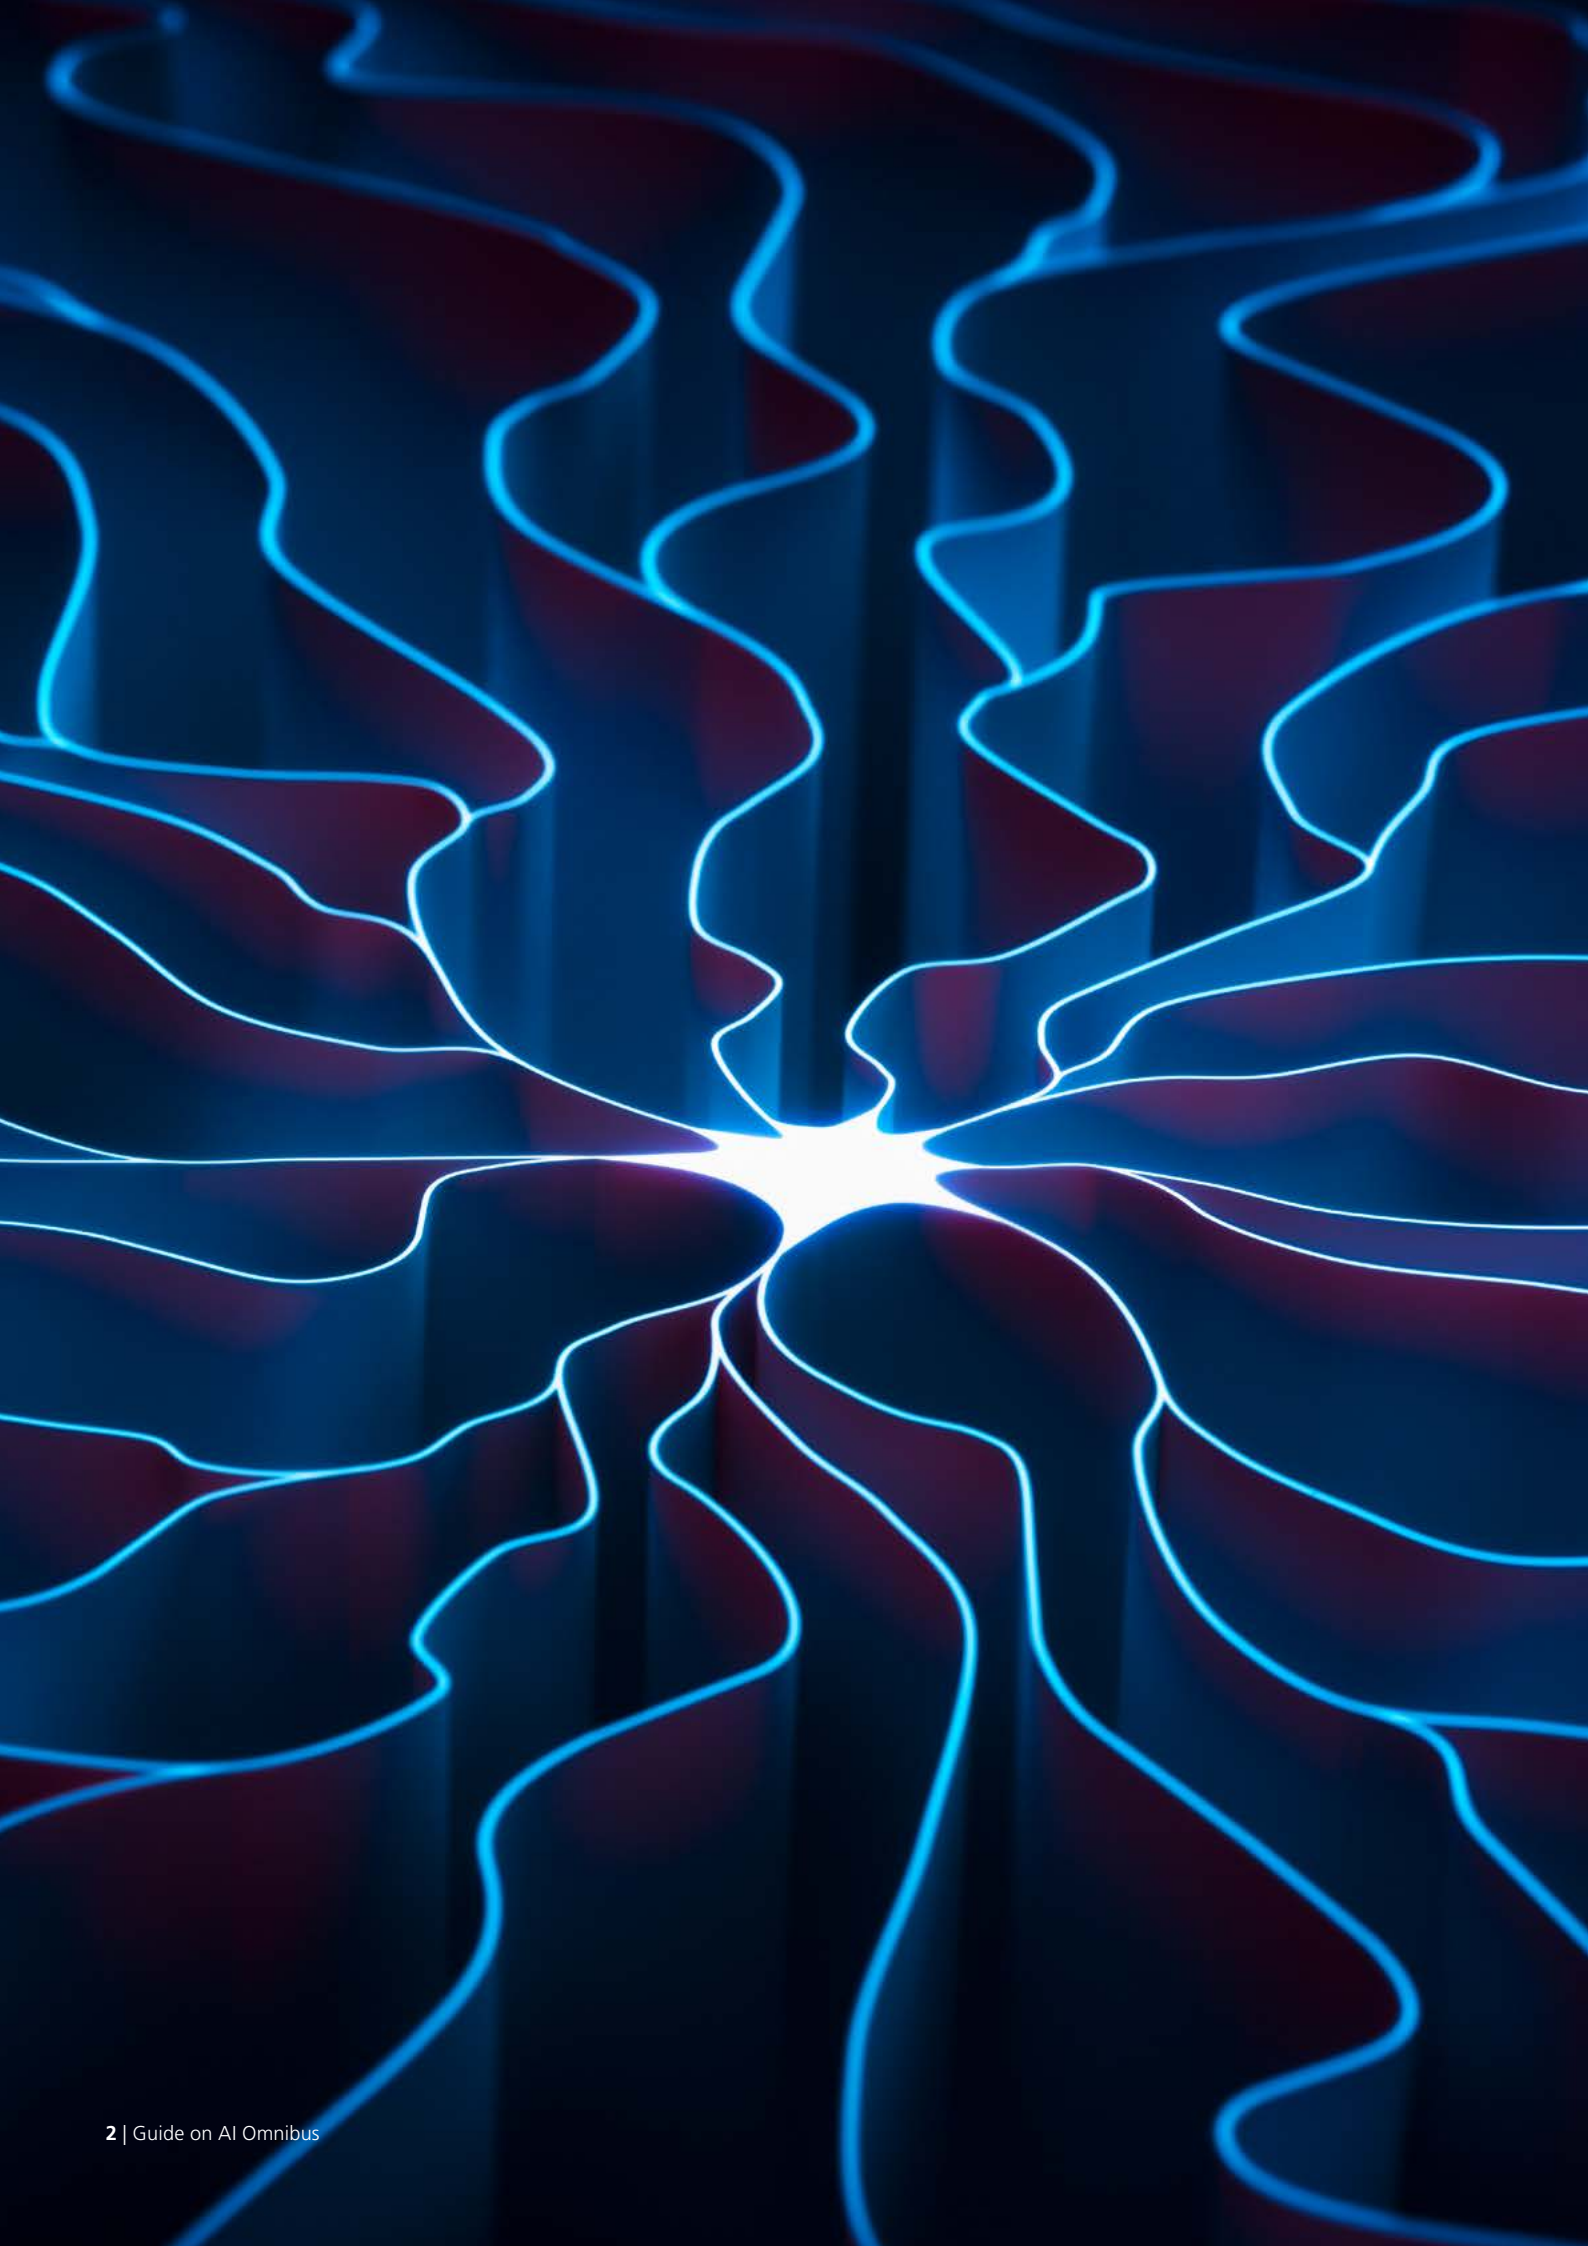

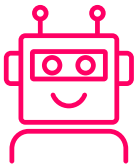

Hello! I am your AI Omnibus Guide

I hope I will be the first text to help you understand what the Digital Omnibus on AI really means for businesses from an AI perspective. **And – although I talk about AI systems – I was created by humans**.

Who created me?

This Guide is promoted by Microsoft and created by CMS lawyers. You can read more about the authors at the end of this Guide.

What is the AI Omnibus?

The Digital Omnibus on AI[1] - let's call it the "**AI Omnibus**" is an EU Regulation amending the European Union's comprehensive framework for artificial intelligence – the AI Act. The AI Act is one of the most significant pieces of EU legislation since the GDPR, affecting the vast majority of enterprises and organisations that develop or use AI systems in their operations. It introduces a package of targeted changes: extended timelines, clarified definitions, new prohibited practices, simplified compliance pathways and strengthened cooperation duties across the AI value chain. Importantly, the proposal does not fundamentally alter the core architecture of the AI Act. The risk-based classification structure and the main obligations relating to high-risk AI systems, general-purpose AI models, and rules on prohibited practices remain generally unchanged.

What is this Guide for?

This Guide helps you navigate the most important changes introduced by the AI Omnibus. It explains them in clear, practical terms to help businesses, both providers and deployers of AI systems, understand what is coming and prepare accordingly. **The Guide is strictly educational**. References to Microsoft's practical experience are included solely to illustrate how the amended AI Act may operate in practice. These references do not constitute product endorsements or commercial recommendations.

1 REGULATION (EU) 2026/1744 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)

Before we jump into the details, let me explain:

What are the key terms and roles you need to know before reading?

Key Terms

AI Act	Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, available here.
AI Omnibus	the EU Regulation 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending the AI Act. It was adopted by the European Parliament on 16 June 2026, endorsed by the Council on 29 June 2026, and published in the Official Journal of the European Union on 24 July 2026. The final text is available here.
AI system	a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the inputs it receives, how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments. *In simple terms: if you use Microsoft Copilot to draft a document, summarize data or generate an image - you are using an "AI system".*
High-risk AI (HRAI)	an AI system that falls under Article 6 of the AI Act.

General-purpose AI model (GPAI model)

an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications. This does not include models used solely for research, development or prototyping before they are placed on the market.

In simpler words, a GPAI model is a broadly capable AI "engine" that can be integrated into or form part of an AI system and be used for many different purposes.

General-purpose AI system (GPAI system)

an AI system based on a general-purpose AI model that can serve a variety of purposes, either for direct use or for integration into other AI systems.

For example, Microsoft Copilot is a GPAI system built on a GPAI model (such as GPT-4o). It can be used directly by end-users (e.g. for writing, summarising or coding) and can also be embedded into business workflows, for example via Copilot Studio to create specialised AI agents.

How to differentiate between a model and a system?

AI models are essential components of AI systems, but they are not AI systems in themselves. To become an AI system, a model must be combined with additional elements such as a user interface or integrations.

Think of it this way: the model is the engine, while the AI system is the car you actually drive. For example, OpenAI's GPT-4o is a GPAI model that powers Microsoft Copilot as:

— the GPAI model (GPT-4o) provides the underlying capability, while
— the GPAI system (e.g. Microsoft Copilot) is the user-facing solution built on top of it.

Key Roles

Provider

a natural or legal person, public authority, agency or other body that develops an AI system or a GPAI model or that has an AI system or a GPAI model developed and places it on the market or puts it into service under its own name or trademark – whether for payment or free of charge.

For example, Microsoft is the provider of Copilot and Azure OpenAI Service.

Deployer

a natural or legal person, public authority, agency or other body using an AI system under its authority, except where the AI system is used in the course of a personal non-professional activity.

In practice, if your company decides to roll out Microsoft Copilot for your employees' daily work – it acts as a deployer.

Operator

an umbrella term covering providers, product manufacturers, deployers, authorised representatives, importers and distributors of AI systems or models.

Who is who in the AI value chain?

Looking at Microsoft as an example – it typically acts as the **provider**, while enterprise customers are typically **deployers**. But the roles are not static. A deployer (e.g. your company using Copilot) may, in certain situations, become a provider →

When roles shift?

This may happen, for example, **when** a deployer (e.g. a company using Copilot) changes the intended purpose of a non-high-risk AI system (including a GPAI system) in such a way that it becomes high-risk.

To understand the "intended purpose" of the AI system you are using, review relevant documentation such as instructions, sales or promotional materials, statements and technical documentation.

Copilot's usage scenario in HR

1

Your recruiter uses Copilot purely to improve wording, translate or summarise CVs, but assesses candidates individually, based on its own assessment. There is no role shift = your company is still "just" a deployer which means - less obligations.

2

However, if for example, your recruiting team is using Copilot Studio to build a recruitment agent that scores, ranks and shortlists candidates - your company may potentially become a 'provider' of a high-risk AI system (Annex III covers AI use cases in recruitment).

Why does this matter?

You need to be mindful of **how** you use **AI systems** and assess whether your actions could result in your company becoming a **provider**.

Important disclaimer

This Guide reflects the final text of the AI Omnibus as published in the Official Journal of the European Union on 24 July 2026, following the completion of the EU legislative procedure.

There was time pressure to complete these steps before 2 August 2026 – the date from which the AI Act's high-risk AI obligations would otherwise have applied.

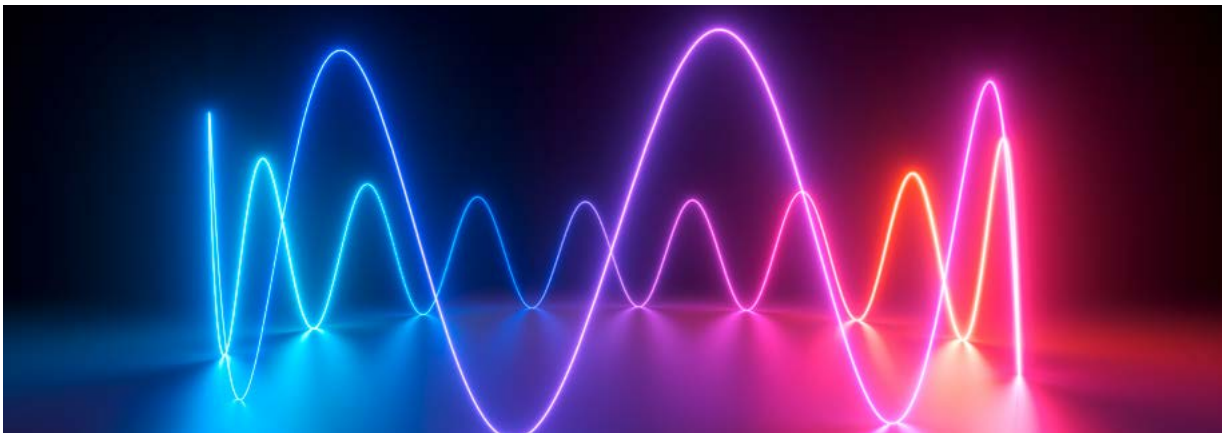

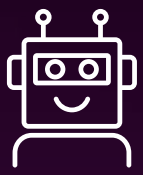

What is inside?

I. Postponed deadlines for High-Risk AI

In a nutshell: As a result of the most significant change of AI Omnibus, obligations for standalone HRAI systems (Annex III) will apply from 2 December 2027 and for HRAI systems embedded in regulated products (Annex I) from 2 August 2028.

What is changing?

The most immediate and tangible impact of the AI Omnibus concerns the timing. The new timeline for compliance has been proposed with key deadlines being pushed back to reflect the current level of regulatory and market readiness. The AI Omnibus proposes for the following:

— **standalone high-risk AI** (listed in Annex III - *e.g. AI systems used in recruitment, credit scoring, education, access to essential private and public services*): **NEW! extended deadline: 2 December 2027** (**formerly**: 2 August 2026)

— **high-risk AI embedded in regulated products** (listed in Annex I - e.g. AI safety components in medical devices): **NEW! extended deadline: 2 August 2028** (**formerly**: 2 August 2027)

What about HRAI already on the market?

If at least one unit of a HRAI was lawfully placed on the market before the relevant new deadlines (see above), other units of the same type and model may continue to be placed on the market, made available, or put into service without additional obligations or requirements, **as long as the design remains unchanged**.

In practice, this means that placing even a single compliant unit of HRAI on the market before the relevant new deadline allows your organisation to continue distributing the same system under the existing regime - without the need to comply with Chapter III obligations (Sections 1–3) of the AI Act, namely HRAI-related rules on classification, technical requirements (risk management, data governance, documentation, transparency, human oversight, accuracy & cybersecurity) and organisational obligations of providers and deployers (quality management system, conformity assessment, registration and FRIA).

This is not only a legal mechanism, but also a strategically valuable opportunity to secure additional time for scaling, product roll-out and compliance planning under the evolving regulatory framework.

Why has AI Omnibus introduced such delays?

— Because harmonised standards and common specifications weren't ready, national authorities are still being set up, conformity assessment bodies hadn't been designated in all Member States.

The result: Keeping the original deadlines would have resulted in disproportionate compliance expectations without the necessary infrastructure to support them.

What does this mean for your business?

— **More time – but not to relax**: This change gives the extra runway for HRAI classification, governance, documentation and conformity assessment planning, but should not be treated as an opportunity to delay compliance efforts.

— **Extended innovation window**: The cut-off date before which high-risk AI systems can be placed on the market outside the Chapter III regime (HRAI classification, technical requirements and provider/ deployer obligations) is also pushed back – giving organisations more space to develop and launch AI solutions before **those** compliance obligations apply.

What should you do now?

— **Use the extended timeline strategically**: Treat the new deadlines (Dec 2027 / Aug 2028) as fixed implementation milestones and use the additional time to build scalable AI governance and compliance frameworks to avoid last minute implementation pressure.

Microsoft Copilot lens:

Microsoft Copilot used for everyday productivity (drafting, summarising, translating) is not a high-risk AI system. However, if your organisation uses for example Copilot Studio to build an AI agent that scores and ranks job candidates, that agent may become a high-risk AI system under Annex III - and then you should observe whether new December 2027 deadline will apply.

What if the AI Omnibus had not been adopted in time? If the AI Omnibus had not been formally adopted before 2 August 2026, the timelines would have remained unchanged, and the AI Act Chapter III obligations would have applied to high-risk AI systems as originally planned, namely from 2 August 2026 for standalone HRAI (Annex III) and on 2 August 2027 for HRAI embedded in regulated products (Annex I).

Guidelines alert: In May 2026, the EU Commission published draft Guidelines on the classification of high-risk AI systems under Article 6. See here. This is a source of practical examples of which AI systems qualify (and which do not) as high-risk. You can't miss it while conducing classification exercise!

II. Transparency: marking and detection for GenAI

In a nutshell: Providers of generative AI systems - already on the market before 2 August 2026 - have a grace period of four months until 2 December 2026 to comply with technical marking and detection requirements relating to AI-generated or manipulated content, while all other transparency requirements apply as originally scheduled from 2 August 2026.

What is changing?

As a rule, transparency obligations under the AI Act will start to apply from **2 August 2026**. The AI Omnibus introduces just **one exception**. It concerns **machine marking and detection obligations for providers of generative AI systems**, where a **four-month delay (until 2 December 2026)** has been proposed. This additional time is available solely for AI systems placed on the market before 2 August 2026. The scope of the postponement is intentionally narrow. It does not cover systems placed on the market on or after 2 August 2026 - those must comply with Article 50(2) from the outset.

Most importantly, the new grace period concerns **only Article 50(2) AI Act**.

All other transparency obligations - including those relating to informing individuals that they are interacting with AI (Article 50(1) of AI Act) and labelling deepfakes (Article 50(4) of AI Act) - *apply in full as of 2 August 2026, without any grace period.*

What does Article 50(2) AI Act actually require?

Article 50(2) imposes a combined obligation on providers, requiring them to ensure both that AI-generated or manipulated content is: (1) **machine-readably marked**, and that it is (2) **detectable** as such. **These two elements are closely interlinked and must be implemented together**. Marking alone is not sufficient without detection, and vice versa. In practice, providers must either offer their own detection tools or rely on interoperable third-party solutions capable of identifying AI-generated or AI-manipulated content. This technical solution must meet a number of qualitative criteria: it must be **effective, reliable, robust and interoperable**.

Guidelines alert: In May 2026 the European Commission published draft Guidelines *on the implementation of the transparency obligations for certain AI systems under Article 50 of Regulation (EU) 2024/1689*, providing practical instructions on marking techniques, detection methods, exceptions and enforcement. The final Guidelines were approved and published on 20 July 2026 - you can read them here.

The Guidelines acknowledge, among others, that no single technological approach currently satisfies all requirements across all types of content. As a result, providers are expected to adopt a combination of techniques - such as watermarking, metadata tagging, cryptographic provenance methods, logging or fingerprinting - to the extent that this is technically feasible.

Scope: To what content AI-marking and detection obligations apply?

To systems generating or manipulating **text, image, audio and video content**, including not only classic generative AI systems but also **GPAI, agentic systems and digital twins** that provide a virtual replica of persons, physical objects or systems.

The key question is: **whether the AI substantively alters the content or its semantics?** Where such a change occurs, marking is required. The Guidelines confirm that this will typically include situations where AI system generates new content (such as translations), inserts new elements into existing content, alters appearance or identity-related features, or creates composite or synthetic outputs. By contrast, purely technical or minor modifications - such as spell/grammar checking, noise reduction, minor cropping or rescaling - do not trigger the Article 50(2) obligation.

No retroactivity: The marking and detection obligations do not apply retroactively. AI-generated or manipulated content that was already made available on the market before 2 August 2026 does not need to be marked or labelled. However, voluntary labelling of such content is encouraged in light of the transparency objectives pursued by the AI Act.

What does this mean in practice?

For providers: taking necessary steps to comply with Article 50(2) by designing combined marking and detection solutions, testing interoperability and aligning with emerging Guidance and the Code of Practice. The additional 4-month transition period should be used strategically and only towards AI systems placed on the market before 2 August 2026.

In late 2025 the European Commission began work on a *Code of Practice on Article 50 compliance*, focusing on Article 50(2) (marking and detection) and Article 50(4) (labelling). Adherence may help businesses demonstrate compliance. **On 10 June 2026, the Commission published the final and approved version of the Code of Practice**.

For deployers: watermarking obligations do not apply directly. However, deployers should ensure contractually that the solutions they rely on produce outputs that are properly marked and detectable, so that their own obligations can be fulfilled on a compliant basis.

Microsoft Copilot lens

Role-shift scenario: If your organisation uses Azure OpenAI models to build its own AI system and puts it into service under its own name, it may qualify as a provider. In such a case, the marking and detection obligation under Article 50(2) will apply directly. Microsoft's upstream compliance does not exempt you from your own provider obligations.

What if the AI Omnibus had not been adopted in time? In that scenario, Article 50(2) obligations would have applied in accordance with the original timeline (from 2 August 2026), without any additional transition period.

III. Exempt systems must still be registered - but under simplified rules

In a nutshell: AI Omnibus retains the obligation to register certain exempt AI systems listed in Annex III in the EU HRAI database, while simplifying the scope of information that must be disclosed.

What is changing?

Under Article 6(3) of the AI Act, some systems that formally fall within HRAI use cases covered by Annex III (including e.g. AI used in recruitment, education, credit scoring or life insurance) may qualify for an exemption and therefore be excluded from HRAI classification. This is typically the case where the system **just performs a narrow procedural task** (e.g. organising received CVs into a searchable database), **improves the result of a previously completed human activity** (e.g. linguistic correction of a draft prepared by a human), detects patterns in prior **decision-making without replacing human judgment** (e.g. flagging anomalies in past lending decisions for a human reviewer), **or performs a preparatory task to an assessment** (e.g. pre-sorting applications before a human evaluator takes over) - provided in each case that the system does not pose a significant risk of harm to health, safety or fundamental rights.

Example from the Guidelines on High-Risk AI Classification: Guidelines provide examples of AI systems used in recruitment that simply recognise and organise information from received CVs and store them in a searchable database for recruiters. While such systems formally fall within a high-risk use case in Annex III, they may benefit from Article 6(3) exemption, where they only perform a clearly defined and limited function and do not materially influence hiring decisions, as they are considered to support merely preparatory or procedural tasks.

Despite initial discussions on potentially removing the registration obligation for exempt systems altogether, the **AI Omnibus retains it while simplifying the process.**

What does this mean in practice?

The AI Omnibus keeps the registration obligation for exempt high-risk AI systems while simplifying the overall process. In particular, fewer data fields must now be completed when registering these systems in the EU HRAI database. **Notably, providers are no longer required to: publicly disclose their reasoning for why the system qualifies for a HRAI-exemption or indicate in which Member States it is made available.**

The remaining registration fields cover only core identification elements, such as the provider's details, the system's name and intended purpose, the conditions under which it qualifies for the exemption, and its market status.

Important caveat: the simplification concerns only the scope of information that must be disclosed in the EU HRAI database. It does not reduce the provider's underlying compliance obligations. Providers must still carry out and document their HRAI-exemption assessment prior to placing the AI system on the market, including that the system does not pose a significant risk to health, safety or fundamental rights. This documentation may still be requested at any time by competent authorities.

In simple terms, the assessment must still exist in full - it simply no longer needs to be publicly disclosed.

For businesses, the key takeaway is that exemption does not eliminate compliance steps. Even systems that are "non-high risk in practice" still trigger formal obligations. Providers must: (1) document their classification assessment under Article 6(3), (2) register the system in the EU HRAI database (although following AI Omnibus - in a simplified form), and (3) be prepared to provide this assessment to authorities upon request. Internal classification decisions must therefore remain well-documented even where external disclosure requirements are reduced.

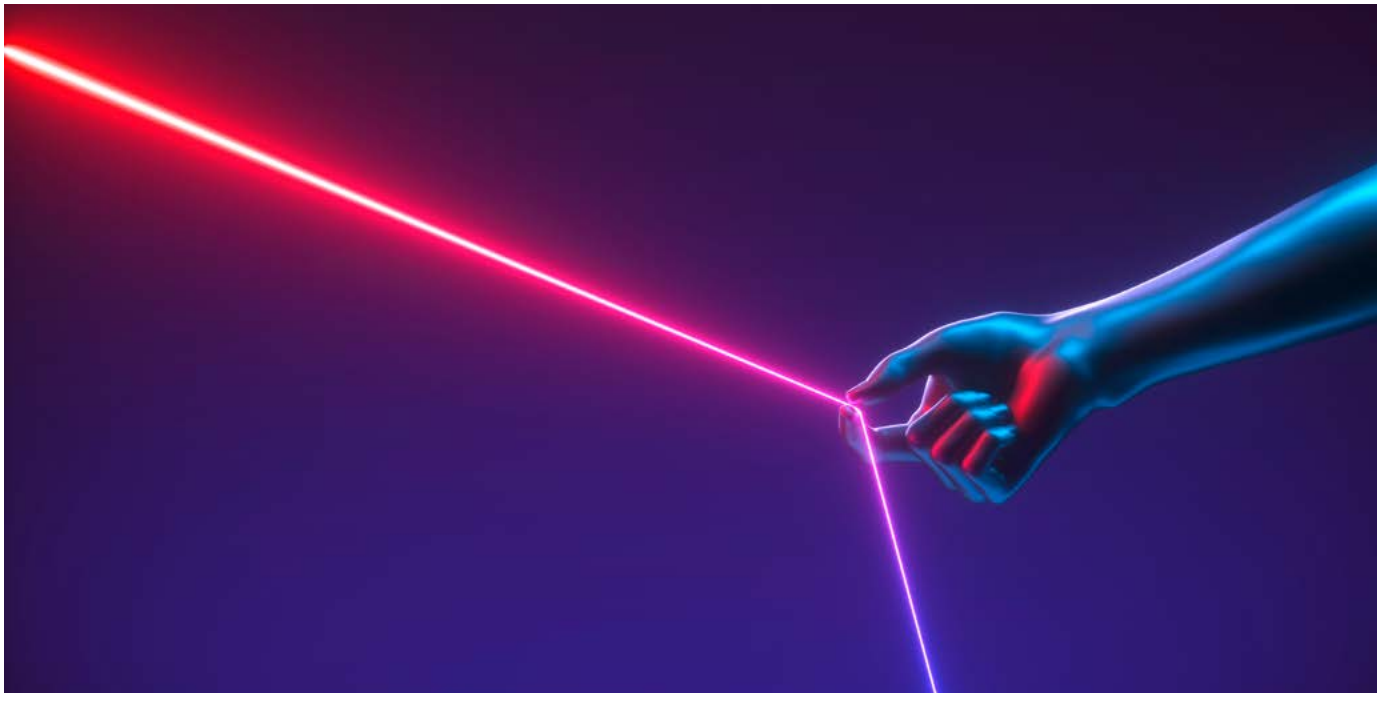

IV. New Prohibited Practices: nudifiers

In a nutshell: The AI Omnibus adds two new prohibited practices to Article 5 of the AI Act, targeting AI-generated non-consensual intimate material (**NCIM**) and child sexual abuse material (**CSAM**). Both prohibitions will apply from 2 December 2026, and address risks that could not have been fully anticipated when the AI Act was originally drafted.

What is changing?

The list of prohibited practices in Article 5 of the AI Act was never intended to be static. It is subject to ongoing review by the European Commission as technology evolves. The AI Omnibus marks the first exercise of that review. It illustrates what we describe as "unknown unknowns": risks that legislators could not fully anticipate when the AI Act was originally drafted, as they emerge from rapidly advancing AI capabilities.

As a response to this, AI Omnibus **adds two new prohibited practices to Article 5, applying from 2 December 2026:**

Non-consensual intimate material (NCIM)

The first prohibition targets AI systems that generate or manipulate realistic images, videos, audio or similar material depicting an identifiable natural person's intimate parts or sexually explicit activity - without that person's explicit consent.

*The term **realistic** refers to a credible depiction of the person's face, voice or body in a way that resembles real-life appearance, even if it does not fully correspond to the person's actual features.*

This **does not** cover clearly unrealistic or cartoonish depictions, or legitimate uses such as certain medical imaging (e.g. mammograms).

Child sexual abuse material (CSAM)

The second prohibition targets the placing on the market, putting into service, or use of an AI system that generates or manipulates material or performance within the meaning of Article 2(c) and (e) of Directive 2011/93/EU - namely child pornography and pornographic performances involving children - including wholly or partially synthetic material.

A "child" is any person below the age of 18.

Critically, the Directive's definitions go beyond depictions of real children. They also capture material depicting any person appearing to be a child and wholly synthetic realistic imagery.

This means that the argument that "no real child was harmed" will not be a valid defence.

New prohibitions – what triggers?

For providers: the ban applies not only when the system is intentionally designed to generate prohibited material, but also where such generation is a reasonably foreseeable and reproducible outcome and the provider has failed to implement adequate safeguards.

What does "adequate safeguards" mean in practice? AI Omnibus offers a helpful checklist. Think of it as a layered defence: clean your training data so the model doesn't learn from prohibited material, train it to refuse harmful prompts, put guardrails around what users can ask and what the system can output, monitor for patterns of misuse, and give users a way to report abuse. Providers should document safeguards and abuse-detection measures before 2 December 2026 asking two control questions: whether the system can produce prohibited content, and whether foreseeable misuse is adequately prevented?

For deployers: the ban targets intentional use of the system to generate prohibited material - including cases where users actively work around the provider's safeguards. Accidental generation should not trigger liability, but in practice, the circumstances may still influence this assessment.

This is exactly why AI governance matters. If your organisation controls and monitors AI use and applies clear rules, you put yourself in a much stronger position if something goes wrong. Acceptable-use policies, approved-tool governance and employee awareness of the AI tool proper use are the key of AI Governance. When employees rely on unapproved AI tools (shadow AI), the organisation's exposure to prohibited-practice risk increases - because it cannot demonstrate that it governed, monitored or prevented such use.

V. New legal pathway to process sensitive data to detect AI bias

In a nutshell: The AI Omnibus introduces a new, broader legal basis (Article 4a) allowing also providers and deployers of any AI system to process special-category data (e.g. ethnicity, health, sexual orientation) for bias detection and correction (mitigation). This replaces and broadens the narrower rule that previously applied only to providers of high-risk AI systems. Article 4a applies immediately upon entry into application of the AI Omnibus.

What is changing?

AI systems can reinforce discrimination, as biased training data may lead to outcomes that disproportionately affect certain groups, for example in recruitment, credit scoring or access to public services. To detect and correct such bias, organisations often need to analyse how different groups are treated, which in practice requires processing sensitive personal data (Article 9 of GDPR), such as e.g. ethnicity or sexual orientation. Without processing such data, it may be impossible to ensure that the AI system does not discriminate against certain groups.

The difficulty is that the GDPR allows processing of sensitive data only if one of the Article 9(2) bases applies. The original AI Act addressed this only partially - by Article 10(5) of AI Act - which allows providers of high-risk AI systems to process sensitive data for bias testing under strict and limited conditions. AI Omnibus removes this provision and replaces it with a broader Article 4a, which directly connects the AI Act with the GDPR framework.

As a result, the logic becomes clearer. GDPR allows sensitive data processing where for example EU law defines a substantial public interest and sets appropriate safeguards (Article 9(2)(g)). *Article 4a of the AI Act identifies bias detection as a public interest, requires strict necessity, and introduces specific safeguards*.

Ready-to-go legal basis for your privacy documentation

If your organisation decides to process special-category data for AI bias detection, the legal basis chain for your Records of Processing Activities (RoPA) and privacy notices would most likely be:

— **GDPR Article 9(2)(g)** - processing necessary for reasons of substantial public interest, based on Union law;

— **"Union law" = Article 4a of Regulation (EU) 2024/1689 (as amended by the AI Omnibus)** - which establishes bias detection and correction as the substantial public interest, imposes strict necessity as the proportionality threshold, and prescribes the specific safeguards (pseudonymisation, access controls, deletion obligations, no onward transfers, documentation of necessity).

In practice, your privacy documentation should reference both provisions together and demonstrate compliance with all conditions set out in Article 4a.

Quick note: *However, as the EDPS and EDPB flagged in their Joint Opinion of 20 January 2026, the permissive wording of Article 4a (2) ("**may**" instead of "**shall**") could create legal uncertainty as to whether it genuinely constitutes a sufficient Union law basis for the purposes of Article 9(2)(g) GDPR.* ***Privacy practitioners should monitor further guidance on this point****.*

Who can rely on Article 4a?

The new rule reflects how bias arises in practice - often at deployment stage, not just during development.

Example: *A credit-scoring model may show no bias in development stage but produce discriminatory outcomes when applied to a particular customer base.*

Extending the legal basis to deployers means they can now lawfully audit AI systems for discrimination using their own operational data - something that was legally uncertain before. **It is therefore a good move to extend the broader legal basis for bias detection and correction beyond high-risk providers to include also:**

— deployers of high-risk AI systems,

— providers and deployers of any other AI systems or models.

Article 4a does not require bias testing. It simply removes a legal (including GDPR) barrier to do so. It is also closely linked to existing duties - in particular, the obligation for deployers of high-risk AI to ensure input data are **sufficiently representative** (Article 26(4)) which in practice may require bias analysis. For providers of high-risk AI systems, Article 4a is closely linked to their existing obligation to detect, prevent and mitigate biases under Article 10(2)(f) and (g) of AI Act.

Conditions and safeguards

Such processing is allowed only if all the following conditions are met: bias detection genuinely cannot be achieved using other data (including synthetic or anonymised data); the data is subject to technical limitations on re-use, state-of-the-art security including pseudonymisation; strict access controls are in place; the data is not transmitted, transferred or accessed by other parties; and the data is deleted once the bias is corrected or the retention period ends, whichever comes first. Records of processing must explain why processing was strictly necessary and why other types of data were insufficient.

In practice this requires all hands-on deck: legal, data protection, and IT teams will need to work together to make practical use of this.

Additional threshold for non-high-risk AI systems

For providers and deployers of AI systems that are not classified as high-risk, Article 4a adds (to above conditions) one additional point: the biases must be likely to affect health and safety, have a negative impact on fundamental rights, or lead to discrimination prohibited under Union law – especially where data outputs influence inputs for future operations. This prevents the provision from being used as a blanket justification for processing sensitive data in low-risk contexts.

Microsoft Copilot lens

If your company uses Copilot Studio to build a custom AI agent and observes discriminatory patterns in its outputs, Article 4a gives you a lawful pathway to audit those patterns using special-category data - but only where synthetic or anonymised data cannot achieve the same insight, and only with all safeguards in place. This is a permission to investigate, not a licence to collect broadly.

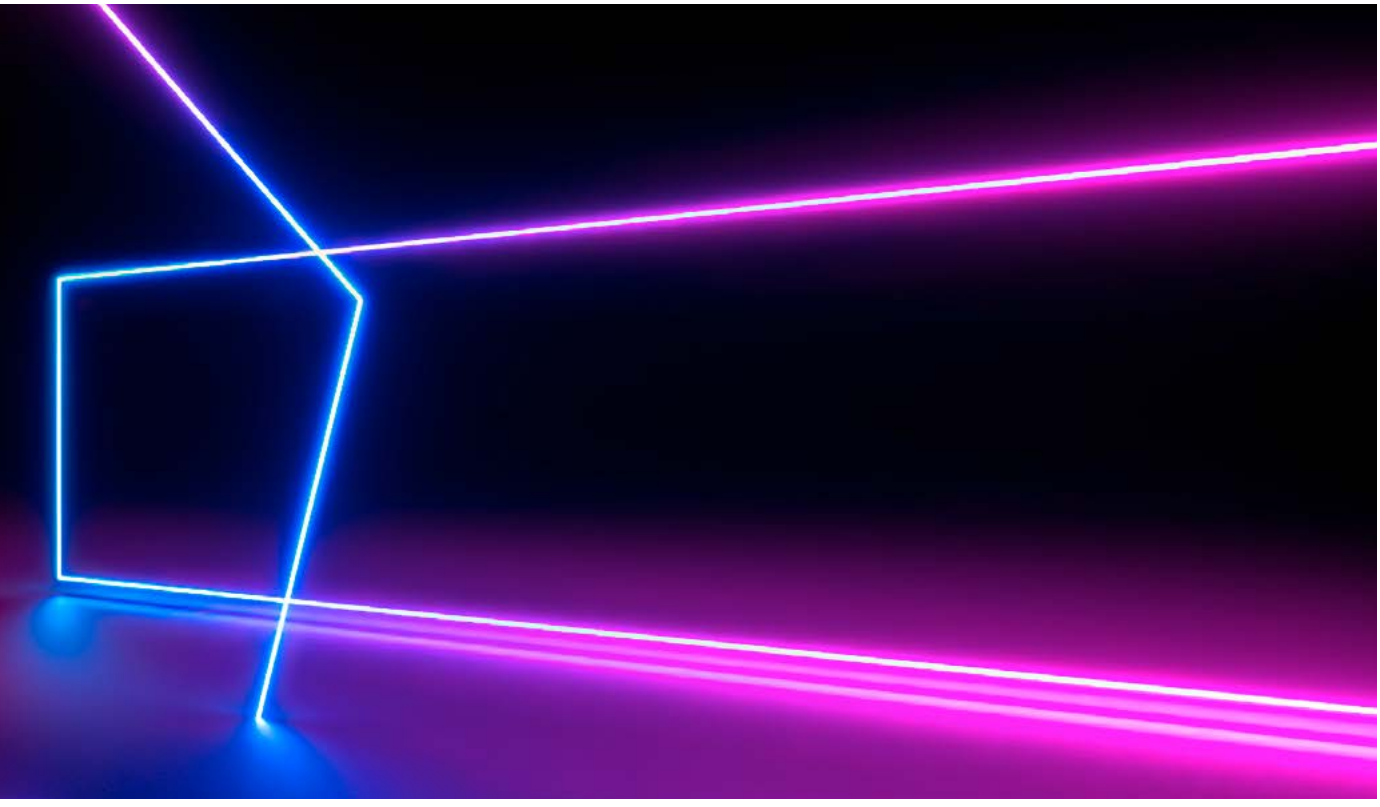

VI. SME and SMC simplification: compliance proportionate to size

In a nutshell: The AI Omnibus extends simplification measures beyond SMEs to small mid-cap companies (SMCs). The key idea is simple: compliance scales to the organisation's size - but the level of protection stays the same. Smaller companies get lighter processes, not lower standards.

What is changing?

The AI Act was originally designed with large technology companies in mind. However, many AI systems are developed and used by smaller organisations - start-ups, scale-ups and mid-sized firms that do not have the same resources or compliance infrastructure. More importantly, most of the Union companies are small and medium-sized enterprises, the majority of which are micro and small enterprises. The AI Omnibus acknowledges this reality. It extends existing simplifications - previously reserved for SMEs only - to a new group: small mid-cap companies (SMCs) and formally embeds both definitions into the AI Act.

The rationale draws on a well-known regulatory phenomenon that SMEs and SMCs play a vital role in the Union's economy. Extending measures to SMCs is just to ease their innovation by avoiding a regulatory "cliff edge" at the 250-employee mark.

Who qualifies?

— An SME (small and medium-sized enterprise) is a company with fewer than 250 employees and annual turnover not exceeding €50 million (or balance sheet not exceeding €43 million), as defined in Commission Recommendation 2003/361/EC. This group includes microenterprises and start-ups.

— An SMC (small mid-cap enterprise) is a company that has outgrown SME status but employs fewer than 750 people and whose annual turnover does not exceed €150 million (or whose balance sheet total does not exceed €129 million), as defined in Commission Recommendation (EU) 2025/1099.

Including SMCs recognises that these companies - while no longer technically "small" - often still operate like SMEs in practice and lack the compliance capacity of big tech players.

What does simplification mean in practice?

The AI Omnibus extends a range of accommodations to both SMEs and SMCs. Importantly, it does not reduce what you need to comply with - but it adjusts how you demonstrate compliance, making the process more proportionate to your organisation's size.

Documentation and quality management. Providers who qualify as SMEs or SMCs won't need to prepare the full Annex IV technical documentation from scratch. The Commission will provide a simplified template tailored to their needs, and notified bodies must accept it in conformity assessments. Similarly, quality management system (QMS) requirements under Article 17 of the AI Act must be implemented proportionately to the provider's organisational size - this applies to both SMEs and SMCs. There is also a fully simplified QMS regime under Article 63, but this one goes a step further and is available only to SMEs (not SMCs), and only where the SME does not have partner or linked enterprises.

Reduced exposure to fines. The AI Act sets three tiers of fines, ranging from €7.5 million to €35 million (or 1–7% of global annual turnover, whichever is higher). The AI Omnibus introduces a safety net for smaller organisations: whichever amount is lower - the fixed euro sum or the percentage of turnover - applies instead. For SMEs, this cap covers all three tiers: prohibited AI practices, breaches of high-risk AI system **requirements, and supplying incorrect** information to authorities. For SMCs, the cap covers the mid- and lower tiers, but the highest tier - for prohibited practices such as deploying banned AI systems - remains uncapped. This is a deliberate signal: even smaller companies face full consequences for the most serious violations. On top of that, Member States must take the economic viability of SMEs and SMCs into account whenever they impose any penalty.

Sandbox access and guidance. Both SMEs and SMCs receive priority access to AI regulatory sandboxes, including the new Union-level sandbox established by the AI Office. National competent authorities are also explicitly encouraged to provide guidance and advice tailored to smaller organisations, considering the recommendations of the AI Board and the Commission.

What does this mean overall?

The message is clear: compliance is still required - but the way you demonstrate it can be lighter, simpler and more proportionate if you are a smaller organisation.

Microsoft lens

This section matters most for Microsoft partners, ISVs and enterprise customers that qualify as SMEs or SMCs and build high-risk AI systems - for example, using Azure OpenAI or Copilot Studio to create AI agents that fall under Annex III (e.g. recruitment, credit scoring). These businesses can benefit from simplified documentation and proportionate QMS. The key takeaway: your obligations remain the same as for any provider of a high-risk AI system - but the AI Omnibus gives you a lighter way to show you meet them.

VII. AI Sandboxes and real-world testing

In a nutshell: The AI Omnibus makes testing more flexible by extending timelines, introducing an EU-level sandbox and expanding real-world testing to new categories of AI systems. At the same time, it gives Member States an extra year to set up national sandboxes, meaning organisations should actively plan how and where to test their AI systems.

What is changing?

Before placing AI systems on the market, organisations often need to test them, either in controlled sandbox environments or under real-world conditions. The original AI Act introduced both mechanisms, but in practice they were quite narrow: sandboxes were slow to materialise, and real-world testing was limited to a specific category of high-risk AI systems. The AI Omnibus addresses this on three fronts: it gives Member States one extra year to build sandboxes, creates a new centralised testing route at EU level, and significantly broadens the scope of real-world testing outside sandboxes.

National sandboxes: one extra year

Member States were originally required to have at least one AI sandbox operational by 2 August 2026. **AI Omnibus extends that deadline to 2 August 2027**. Sandboxes may also be established jointly with other Member States, and the Commission may provide technical support, advice and tools for their establishment and operation. Given the postponed deadline, organisations should not rely solely on national sandboxes in the short term but also consider the alternative testing pathways that AI Omnibus has expanded.

A new EU-level sandbox - Article 57(3a) AI Act

The AI Omnibus enables the AI Office to establish its own AI regulatory sandbox at Union level. This sandbox is designed for AI systems that fall under the AI Office's exclusive supervisory competence under Article 75(1) - primarily AI systems built on general-purpose AI models where the model and the system are developed by the same provider or within the same corporate group, as well as AI systems that constitute or are integrated into very large online platforms or search engines.

In practice, if your organisation develops both a GPAI model and an AI system built on that model within the same corporate group, that system would fall under the AI Office's exclusive competence - and could therefore be tested in the EU-level sandbox.

Certain categories are excluded from this scope, including AI systems embedded in products covered by Annex I, certain biometric systems (Annex III, point 2), and systems used by law enforcement or in the administration of justice. The EU-level sandbox must be implemented in close cooperation with relevant national competent authorities and must provide priority access to SMEs, start-ups and SMCs. Its establishment does not affect Member States' competence to run their own sandboxes for systems under national supervision.

Broader real-world testing options (beyond sandboxes)

One of the most practical changes is the expansion of testing outside sandbox environments. Under the original AI Act (Article 60), real-world testing was limited to standalone high-risk AI systems listed in Annex III - for example, AI used in education, employment or access to essential services. The AI Omnibus keeps that route in place, but opens two new ones:

AI embedded in regulated products (Annex I, Section A) - such as medical devices, toys or vehicles. Under the amended Article 60, **providers** and **prospective providers** of these systems can now test them in real-world conditions at any time before placing them on the market, subject to a testing plan and appropriate safeguards;

Practical example: *Imagine a company building an AI-powered diagnostic tool that qualifies as a medical device under Regulation (EU) 2017/745 (Section A of Annex I). Under the original AI Act, there was no clear route to test such a product-embedded AI system in real hospital settings outside a sandbox. AI Omnibus, by amended Article 60, allows this - the provider can conduct real-world testing with patients and clinicians before placing the system on the market, provided a testing plan is agreed with the competent authority and appropriate safeguards are in place.*

Machinery and other Section B products (new Article 60a) - most notably machinery, which AI Omnibus itself moves from Section A to Section B. These products live in a different regulatory world: they are primarily governed by their own sectoral legislation (in case of machinery - the Machinery Regulation (EU) 2023/1230) and are largely outside the AI Act's direct scope. That said, the sectoral rules will eventually need to incorporate the AI Act's high-risk requirements - things like risk management, data governance, transparency, accuracy and cybersecurity (Articles 8–15). The AI Omnibus lets Member States get ahead of this transition. Member States may but are not obliged to set up national testing frameworks that allow providers to test their AI systems against those high-risk requirements in real-world conditions, even before the sectoral rules formally catch up. If a Member State decides to go this route, it must notify its framework to the Commission beforehand and ensure a high level of protection for health, safety and fundamental rights. The testing must also comply with any applicable provisions of the relevant sectoral legislation itself.

Practical example: *A manufacturer installs an AI safety system in an industrial robot – the system tracks worker movements and adjusts the robot to avoid collisions. This is machinery under the Machinery Regulation, which the AI Omnibus moves to Section B. Before placing it on the market, the manufacturer wants to test the AI in a real factory setting. If the Member State where the factory is located has adopted a testing framework under Article 60a, the manufacturer can run those tests against the AI Act's high-risk requirements. If not, this route is not available - and the manufacturer may need to test in a Member State that has.*

In both cases, sandbox participation is not required – real-world testing can take place independently. The key difference is that the Section A route (amended Article 60) will be available EU-wide by default, while the Section B route (new Article 60a) will only become available where a Member State has actively opted in by adopting a national testing framework.

Sandbox vs. real-world testing — what's the difference?

An AI **regulatory sandbox** is a controlled environment set up by a national authority (or the AI Office at EU level). You enter it before placing your AI system on the market, work under the authority's supervision, and follow an agreed sandbox plan. The sandbox lets you develop, train, test and validate your AI system for a limited time, with appropriate safeguards in place. Think of it as a supervised "safe space" where you can innovate while staying in close contact with the regulator.

Real-world testing is different. It means taking your AI system out of the lab or simulated environment and testing it in actual, real-life conditions - for example, in a hospital, on a factory floor, or in a live recruitment process. The goal is to gather reliable data and verify whether the system meets the AI Act's requirements before you place it on the market. Real-world testing is not the same as placing the system on the market - as long as you meet the conditions set out in the AI Act (such as an approved testing plan and appropriate safeguards), it does not count as making the system commercially available.

Can you combine them? Yes. Real-world testing can take place inside a sandbox (as part of the supervised project) or outside one (independently, under Article 60 or the new Article 60a). If you do both, the AI Omnibus lets you merge the sandbox plan and the real-world testing plan into a single document.

Practical tips

Monitor your national sandboxes' availability. National sandboxes may not be operational until mid-2027 at the earliest, and the EU-level sandbox has no confirmed launch date. In the meantime, the expanded real-world testing routes give you a viable alternative - but you will need a testing plan agreed with the competent authority and appropriate safeguards in place before you start.

Check your Member State's position on Section B testing. The new Article 60a route for products like machinery is not automatic - it only becomes available if your Member State decides to adopt a national testing framework. Since Member States are not obliged to do so, availability will vary across the EU. If you operate cross-border, you may find that testing is permitted in one country but not in another - so check early and plan accordingly.

Combine your plans where possible. If you are participating in a sandbox and planning real-world testing, AI Omnibus allows you to merge both into a single document - a sandbox plan that incorporates the real-world testing plan. This is a small but useful procedural simplification that avoids duplicating documentation and approval processes.

VIII. Streamlined Compliance: AI Literacy and Cybersecurity

In a nutshell: The AI literacy obligation has been softened - from an expectation of outcome to an obligation of effort. At the same time, high-risk AI systems that meet the requirements of the Cyber Resilience Act (CRA) are deemed compliant with the AI Act's cybersecurity requirements, eliminating the need for a separate assessment.

What is changing?

AI Omnibus reduces compliance friction in two areas where the AI Act created disproportionate or duplicative burdens: AI literacy and cybersecurity. In both cases, the core of the obligation is preserved but the way it must be demonstrated is simplified.

AI LITERACY: path from "ensure" to "support"

The original AI Act required providers and deployers to ensure a sufficient level of AI literacy among their staff. The AI Omnibus not only softens the approach - from "ensure" to "take measures to support the development of" - but also clarifies that this does not require guaranteeing a specific level of knowledge. The focus moves from proving an outcome to demonstrating reasonable efforts taken.

Is AI literacy just about Article 4?

AI literacy is not just about Article 4. It is a broader concept that also underpins other AI Act obligations. Especially those relating to high-risk AI systems, which remain untouched by the AI Omnibus.

— Including Article 26(2) of the AI Act, **deployers** must assign human oversight to individuals who have the *necessary competence, training and authority, as well as the necessary support - AI literacy.*

— **Providers**, in turn, must design their systems so that effective human oversight is actually possible (Article 14(4) AI Act). These are not soft expectations but rather binding requirements with real teeth. Also, Article 14(4)(b) requires that persons assigned to human oversight remain aware of the risk of automation bias - the tendency to over-rely on AI outputs without questioning them. System design alone cannot prevent this. It also takes people who understand what the AI can and cannot do and who are trained to challenge its results.

Thus, although the general AI literacy bar (Article 4) has been lowered by the AI Omnibus, the high-risk-specific competence bar remains intact.

Make AI Literacy your priority

Recital 8 of the AI Omnibus makes this explicit:

"*AI literacy should be a strategic priority, regardless of regulatory obligations and potential sanctions.*"

If an incident occurs and staff lack the basic understanding to use AI systems safely, the **absence of structured literacy efforts, as a part of your AI Governance, will be difficult to justify.**

TIP - Shadow AI: Article 4 does not address this directly, but organisations should manage the growing risk of employees using unapproved AI tools outside formal processes. The key are clear internal policies on which tools may be used, visibility over AI usage across the organisation, and where necessary the companies should consider prohibition of shadow-AI (non-authorised) tools for work purposes.

Institutional support. The AI Omnibus adds a support layer. Under the new Article 4(2), the Commission and Member States must help providers and deployers - in particular SMEs - by publishing practical compliance examples and adopting common objectives through the AI Board. The Commission is also developing a repository of AI literacy practices that can serve as a practical reference point.

CYBERSECURITY: one assessment instead of two

High-risk AI systems must meet cybersecurity requirements under two separate regimes that largely overlap:

— Article 15 of the AI Act requires providers to ensure an appropriate level of accuracy, robustness and cybersecurity throughout the system's lifecycle - including resilience against AI-specific threats such as data poisoning and adversarial attacks.

— The Cyber Resilience Act (CRA) imposes essential cybersecurity requirements on products with digital elements - including secure-by-design principles, vulnerability handling, and risk assessment that must account for AI-specific risks.

Until now, a provider whose system fell under both regimes had to demonstrate compliance with each one separately - even where the requirements covered the same ground. The original AI Act did offer a limited shortcut: under Article 42(2), a system that obtained a **voluntary** cybersecurity certificate under the Cybersecurity Act (Regulation (EU) 2019/881) could benefit from a **presumption** of compliance with Article 15 of AI Act but since that certification programme is voluntary, few providers used it in practice.

AI Omnibus makes this explicit in the AI Act

The deemed conformity mechanism already exists in Article 12(1) of the CRA itself. AI Omnibus mirrors it in the AI Act by proposing a new Article 42(3) so that providers can see the link directly, without having to look it up in the CRA. Under this provision, if your high-risk AI system falls within the scope of the CRA and meets the conditions laid down in Article 12(1), the system is deemed to comply with the AI Act's cybersecurity requirements under Article 15 - no separate AI Act cybersecurity assessment needed. **This is a significant step up in clarity from the original Article 42(2), which relied on a voluntary cybersecurity certification scheme that few providers used in practice.**

One important catch: this works only to the extent that your CRA declaration of conformity covers the same ground as Article 15 AI Act. Article 15 also addresses accuracy and robustness which the CRA does not cover - so those requirements still need to be met separately.

What does this mean in practice?

For providers, cybersecurity compliance can now be streamlined. Where your system falls within the scope of the CRA and meets its requirements, those efforts can be relied upon for AI Act compliance - avoiding duplication of assessments and documentation. To be clear: this simplification applies only to cybersecurity. All other high-risk AI obligations - risk management, documentation, human oversight, transparency - remain fully in force.

Practical tip:

Map your CRA Article 12(1) requirements against AI Act Article 15 line by line. Where they overlap, one assessment is enough. Where they don't - fill the gaps independently.

What if the AI Omnibus had not been adopted? On AI literacy, the obligation would have remained in its original, stricter form - requiring organisations to ensure a sufficient level of AI literacy rather than merely take measures to support it. On cybersecurity, the deemed conformity mechanism would have still existed under Article 12(1) of the CRA itself - but it would have not been reflected in the AI Act, making the interplay between the two regimes less visible.

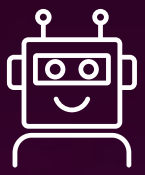

Sector–related changes

Who should read this?

Read it especially if you represent a financial sector or manufacture/use AI-embedded products covered by harmonised sectoral legislation (automotive, machinery, medical devices)

I. Reduced overlap between the AI Act and existing sectoral product frameworks

In a nutshell: AI Omnibus aims to reduce overlap between the AI Act and existing EU product safety laws, but the changes are more targeted than expected, mainly affecting **machinery**, while for **other regulated products** offering only more limited simplification measures.

What does simplification mean in practice?

The AI Omnibus focuses on clarifying how the AI Act interacts with existing EU product safety frameworks and reducing duplication for AI embedded in regulated products. However, the changes are more limited than initially expected. The key structural shift concerns machinery only: the Machinery Regulation (EU) 2023/1230 has been moved from Section A to Section B of Annex I to the AI Act.

Section B products: What does it mean for machines?

— **Limited direct AI Act application**: High-risk AI systems used in machinery will no longer be subject to the full AI Act framework. Only selected provisions will continue to apply directly: HRAI classification (Art. 6(1)), real-world testing (Art. 60a), the AI Office enforcement framework (Articles 102–112), and conditionally AI sandboxes (Articles 57–59) once HRAI requirements are integrated into the Machinery Regulation.

— **AI Act requirements shift into sectoral law**: Core AI Act requirements - risk management, data governance, transparency, human oversight, technical documentation, quality management, and post-market monitoring - will be integrated into the Machinery Regulation **via delegated acts (due by 2 August 2028)**. Until sector-specific standards are developed, compliance with AI Act harmonised standards will give a presumption of conformity.

Practical tip:

treat AI Act's standards as a transitional compliance bridge until sectoral measures are adopted.

The AI Act's requirements are not removed but will be absorbed into the Machinery Regulation - resulting in a single compliance framework.

Section A products: still under the full AI Act (but with targeted simplifications)

Against industry expectations, **medical devices (MDR/IVDR) and other regulated products** remain in Section A. This means that manufacturers must still comply **both** with:

— the full set of AI Act's HRAI requirements, and
— the applicable sectoral product legislation.

However, the AI Omnibus offers targeted simplification measures for Section A products instead.

What do "simplification measures" mean in practice for Section A products?

— **Possible targeted reduction of duplication**: The Commission may adopt delegated acts (by 2 August 2027) to streamline selected AI Act obligations, but only where sectoral rules already ensure an equivalent level of protection.

— **Single, coordinated designation of notified bodies**: Notified bodies will be able to undergo a single application and coordinated assessment for both the AI Act and sectoral regimes, reducing administrative burden, particularly for groups.

— **One integrated conformity assessment (both regimes covered)**: Providers may follow single sectoral conformity assessment procedure (e.g. under MDR), which covers also AI Act requirements. This results in one process but covering both regimes in full. Where sectoral rules allow self-assessment, this remains possible, provided all AI Act requirements are also addressed (e.g. via harmonised standards or common specifications).

— **Guidance to streamline compliance**: By 1 August 2027, the Commission will publish guidelines on how to align AI Act and sectoral requirements, supporting integration of risk, quality and documentation processes and reducing duplication in line with proportionality and complementarity principles.

Section A vs Section B – key distinction

— **Section A products** *(e.g. medical devices, toys, lifts, radio or pressure equipment)* – HRAI systems face full set of AI Act requirements. AI Omnibus offers a few targeted but limited simplifications for compliance.

— **Section B products** *(e.g. machinery - after the AI Omnibus move)* – HRAI systems face very limited application of the AI Act (see above). The substance of the AI Act's requirements will be added to the sectoral legislation, including the Machinery Regulation through delegated acts (due by August 2028). Until then, AI Act harmonised standards give presumption of conformity.

What is the practical impact? From a provider's perspective, if you make AI-enabled **medical devices**, you face two sets of rules simultaneously (MDR + full AI Act). If you make AI-enabled **machinery**, you follow one regime - the Machinery Regulation - which will absorb the AI Act's substance through delegated acts. One pathway, not two.

Key question for your business: Under which legislation is your product regulated? Check Annex I to find your answer - or ask your product safety team.

What if the AI Omnibus had not been adopted? Machinery would have remained in Section A of Annex I - meaning manufacturers would have faced the full AI Act alongside the Machinery Regulation, with no sectoral absorption of AI requirements. The delegated acts mechanism for limiting AI obligations where sectoral rules provide equivalent protection would not have existed.

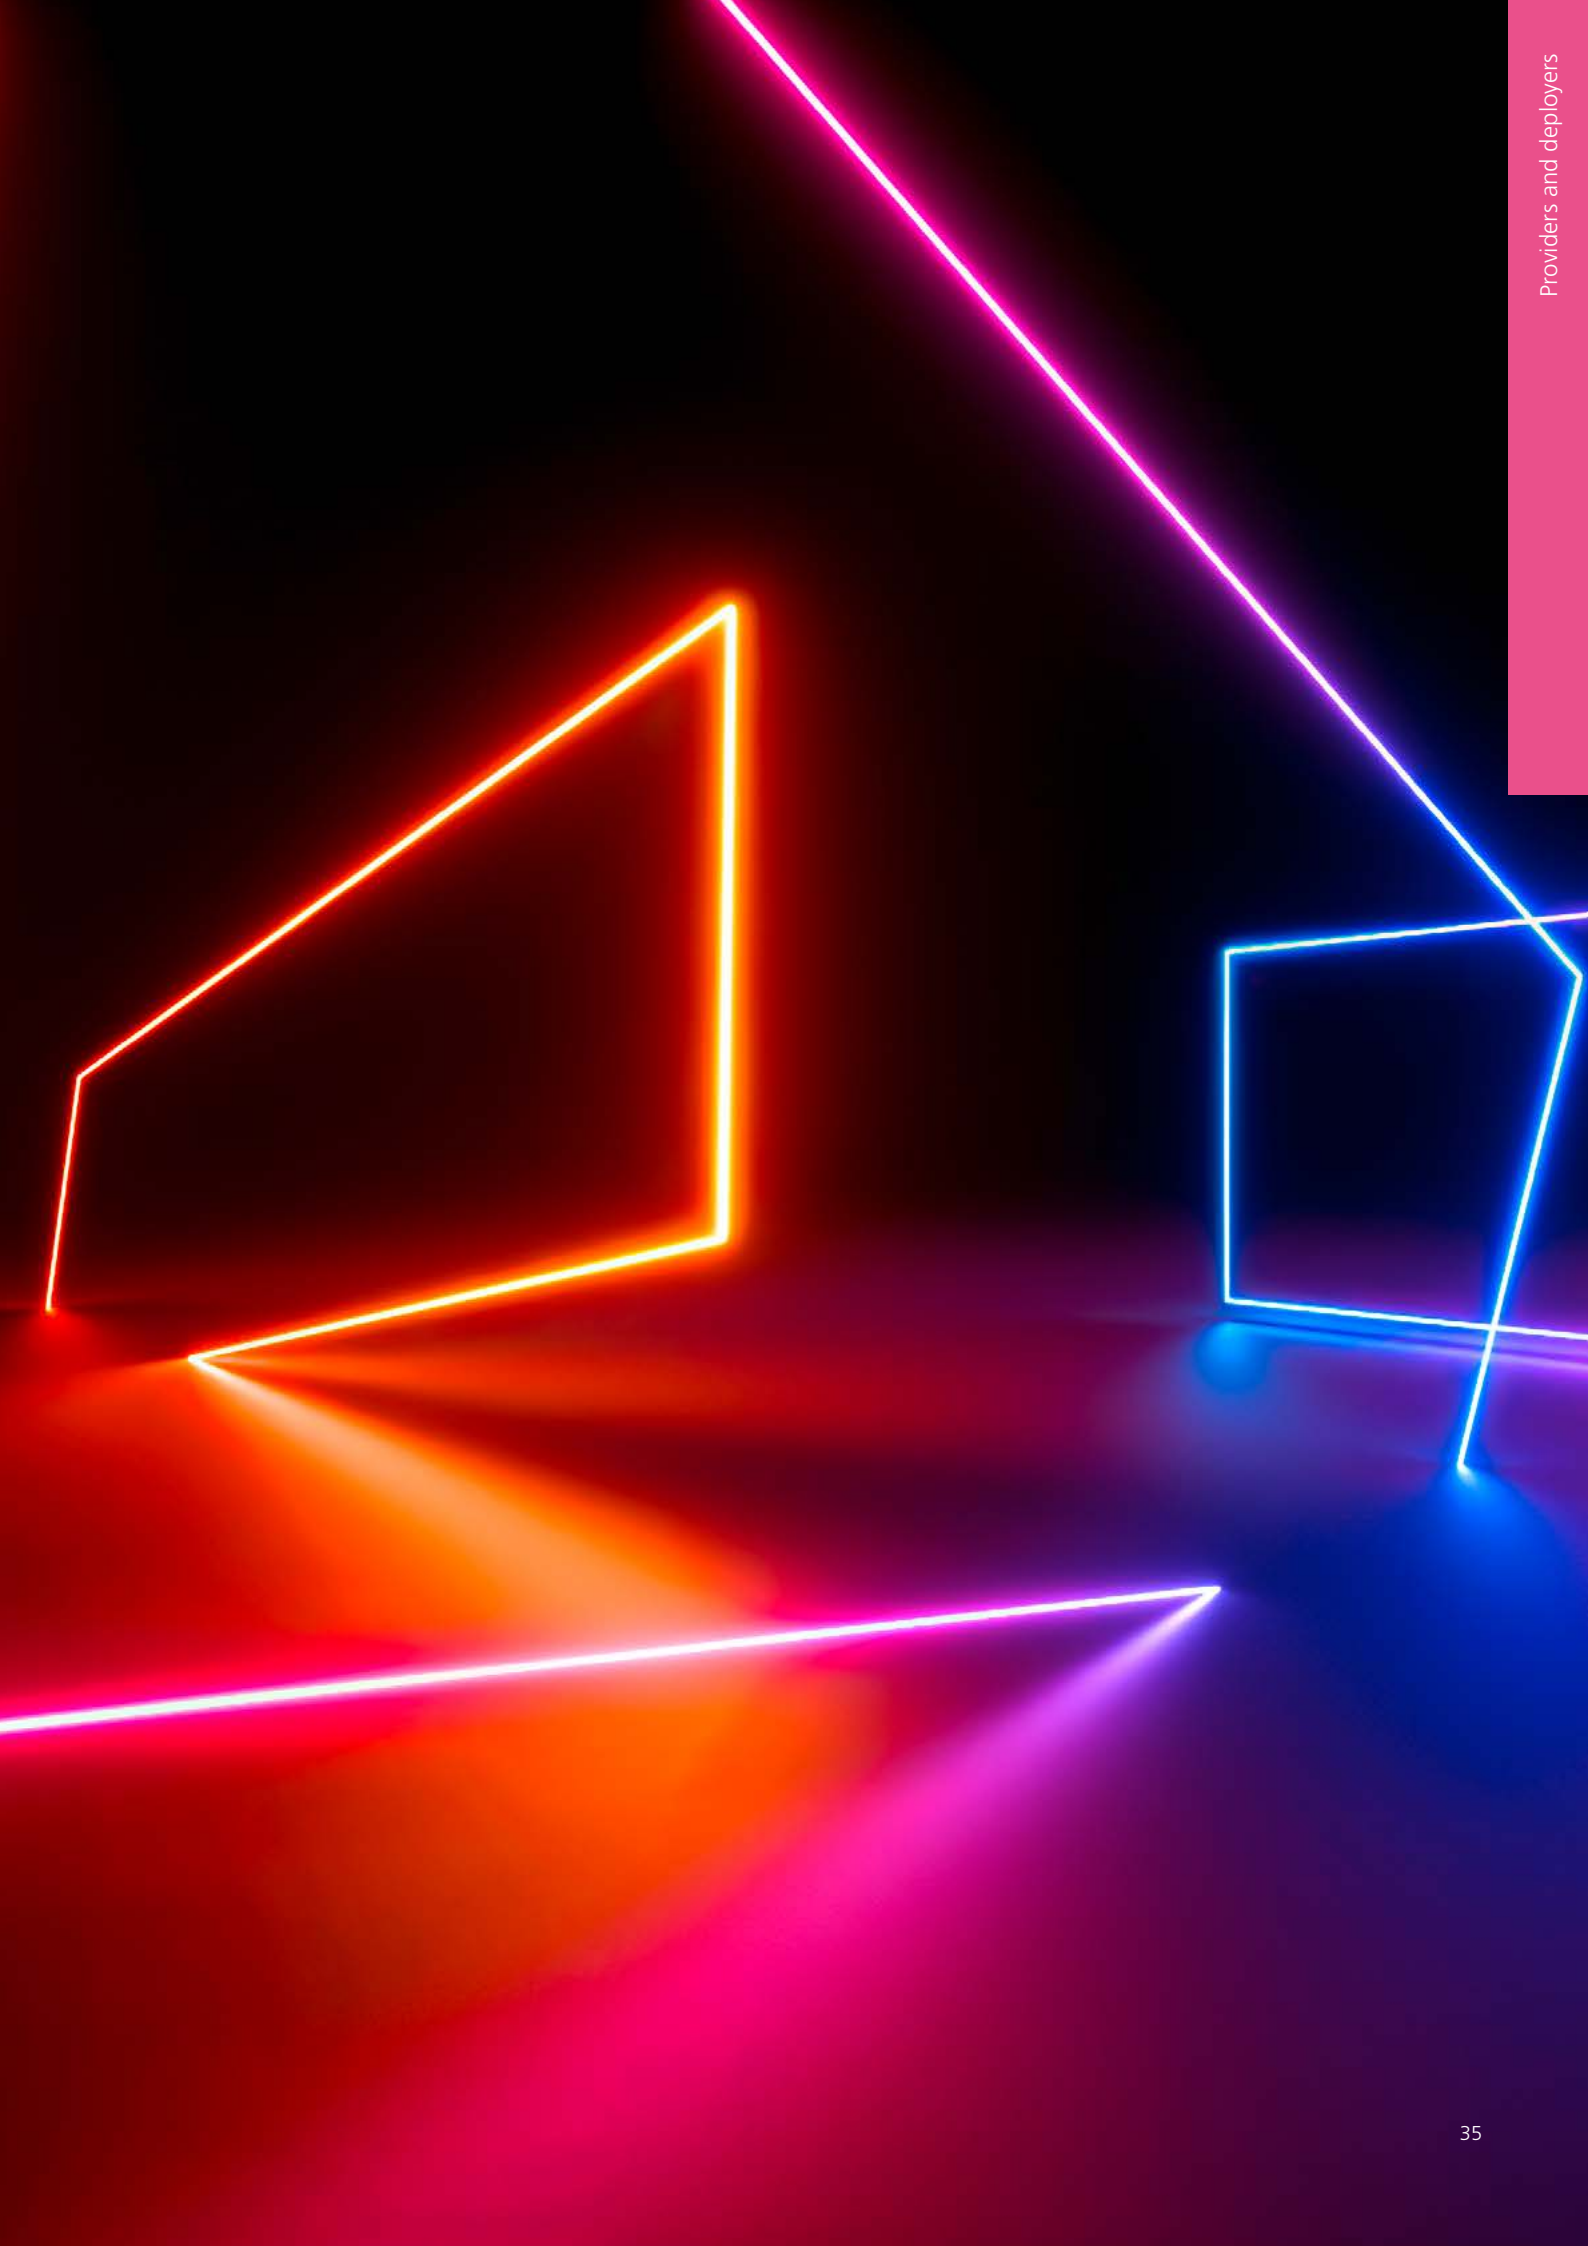

II. FRIA and DPIA: less duplication

In a nutshell: AI Omnibus clarifies the relationship between the fundamental rights impact assessment (FRIA) and the data protection impact assessment (DPIA), allowing deployers to build on existing DPIA work rather than duplicating it.

What is changing?

Brief reminder – what is the current framework under AI Act?

Deployers of HRAI systems in certain categories must carry out a fundamental rights impact assessment (**FRIA**) before putting the system into use. This obligation applies to:

- public bodies and private entities providing public services - for any HRAI system listed in Annex III (except biometric systems under point 2), and
- all deployers (including purely private entities) of credit scoring, creditworthiness assessment, and risk assessment and pricing systems for life and health insurance (banks and insurance companies watch out!).

Where such systems involve processing of personal data, a data protection impact assessment (DPIA) is separately required under the GDPR (Article 35) or the Law Enforcement Directive (Article 27).

In practice, the two assessments overlap significantly - covering similar elements such as risk identification, affected groups and mitigation measures. The original AI Act attempted to address this by stating that the FRIA "*shall complement*" the DPIA (Article 27(4)).

What does AI Omnibus change?

The AI Omnibus replaces the vague "*shall complement*" with a clearer and more practical mechanism: where the relevant FRIA elements are already covered by a DPIA, *the deployer may include cross-references to the relevant DPIA sections or incorporate relevant parts of the DPIA directly into the FRIA* (amended Article 27(4)). Deployers no longer need to reproduce work already done under data protection law - provided that all FRIA requirements are still fully covered.

In addition, the AI Office will develop a template questionnaire - including through an automated tool - to facilitate compliance with the FRIA obligation in a simplified manner (amended Article 27(5)).

Why this matters in practice?

For organisations that already conduct DPIAs under the GDPR - which in practice covers a very wide range of AI use cases involving personal data - this change means they can build their FRIA on top of existing data protection compliance work, rather than constructing a parallel process from scratch. This is particularly valuable for deployers in sectors such as financial services, insurance and public administration, where both assessments are routinely required.

Practical tip:

— **Align internal DPIA and FRIA processes**: design your DPIA methodology so that it also covers FRIA elements (e.g. affected groups, specific harm risks, human oversight, complaint mechanisms) to enable easy cross-referencing.

— **Involve the DPO in AI compliance**: as the DPIA becomes a central building block for FRIA, making early involvement of a data protection officer is essential.

— **Monitor the AI Office's FRIA template**: once published, it will define the structure and scope of cross-referencing. Track its publication timeline.

— **Do not assume full overlap**: DPIA and FRIA serve different purposes - while the DPIA focuses on data protection risks, the FRIA covers broader fundamental rights (non-discrimination, dignity, access to justice). Cross-referencing is allowed, but gaps must be filled and additional analysis may still be required.

Timing note: FRIA obligation sits within Chapter III of the AI Act, which falls within the delayed HRAI timelines - following AI Omnibus: 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems.

Microsoft lens

This is relevant for enterprise customers deploying high-risk AI systems (e.g. via Azure AI or Copilot Studio) in credit scoring, insurance or public services. Where such deployments involve personal data - which is nearly always the case - the ability to build FRIA on top of an existing DPIA reduces the compliance workload significantly.

One last thing and goodbye

At the beginning of this Guide, we said that the AI Omnibus is not a revolution, but a recalibration. Having walked through it, section by section, that assessment holds. The AI Act remains one of the most ambitious pieces of technology regulation in the world. The AI Omnibus does not change that ambition. It adjusts the pace, removes unnecessary friction, and provides businesses - especially smaller ones – with a more realistic path to compliance.

But recalibration is not relaxation. The obligations are real, the deadlines are approaching, and the enforcement architecture is being built as you read this. The organisations that treat the AI Omnibus as breathing room to prepare will be in a strong position. Those that treat it as permission to wait will not.

The AI Act is here. The AI Omnibus is shaping how it will work. Now it is your turn to act.

Key Takeaways

More time - but not a pause

The new deadlines for high-risk AI systems (2 December 2027 for Annex III; 2 August 2028 for Annex I) give organisations additional time to prepare - but not a reason to wait. A compliance framework – which includes system inventories, risk classification, supplier due diligence - takes time to build properly. Treat this as "last call" to start or "signal" to continue building AI governance now! When the deadlines arrive, regulators expect to see mature, documented processes, and not last-minute scrambles.

Transparency kicks soon

If you run chatbots, generate content at scale, or publish AI-assisted materials, the rules of Article 50 apply. Providers of systems placed on the market before 2 August 2026 have until 2 December 2026 to comply with the marking and labelling requirements for AI-generated content (Article 50(2)). All other transparency obligations apply from 2 August 2026. This is not a distant deadline.

New prohibitions demand immediate attention

The provisions on non-consensual intimate imagery and CSAM carry some of the highest penalty exposure in the entire AI Act but more fundamentally, they aim to build a safer digital environment. If your organisation operates image-generation or content-editing pipelines, review your safeguards now. These rules apply from 2 December 2026.

Product AI: know your section

If you make AI-enabled products, your compliance pathway depends on where your product sits in Annex I. Machinery moves to Section B - meaning the AI Act's substance will be absorbed into the Machinery Regulation via delegated acts (due by August 2028). Medical devices, toys and other Section A products remain under the full AI Act, with some targeted simplifications. Start mapping the overlaps now rather than waiting for delegated acts to materialise.

SMEs and SMCs: lighter process, same standards

The AI Omnibus extends simplified documentation, proportionate QMS, reduced fine exposure and priority sandbox access to both SMEs and the new category of small mid-cap companies (up to 750 employees, turnover ≤ €150M). Your obligations do not change - but the way you demonstrate compliance becomes more proportionate. Check your eligibility and plan accordingly.

Testing: don't wait for sandboxes

National sandboxes may not be operational until mid-2027. In the meantime, AI Omnibus opens new real-world testing routes for AI embedded in regulated products (Section A, EU-wide) and for machinery (Section B, where Member States opt in). If you need to validate your AI system before placing it on the market, explore your pathways for testing.

AI literacy: softer wording, same expectation

The obligation has shifted from "*ensure a sufficient level*" to "*take measures to support*". AI literacy softens but does not disappear. If something goes wrong and your staff lack the basic understanding to operate AI safely, the absence of any structured literacy programme will be hard to defend. AI Governance including written policy, training records, clear rules on approved tools - these are your minimum.

Cybersecurity: one assessment where CRA applies

If your high-risk AI system already meets the Cyber Resilience Act requirements, AI Omnibus confirms - directly in the AI Act - that this covers the Article 15 cybersecurity box. Map your CRA compliance against Article 15, identify any gaps, and avoid duplicating work.

FRIA and DPIA: build once, use twice

For deployers in financial services, insurance or public administration, AI Omnibus allows cross-referencing between the fundamental rights impact assessment and the data protection impact assessment. Design your DPIA methodology to cover FRIA elements from the start - you will save significant effort when the HRAI timelines arrive.

Final note

This Guide reflects the final text of the AI Omnibus as published in the Official Journal of the European Union following the completion of the EU legislative procedure.

Authors:

Tomasz Koryzma
Partner, Head of IP/TMT
T +48 22 520 84 79
M tomasz.koryzma@cms-cmno.com

Adriana Zdanowicz-Leśniak
Counsel
T +48 22 520 55 42
M adriana.zdanowicz-lesniak@cms-cmno.com

Agnieszka Czopska
Associate
T +48 22 520 84 28
M agnieszka.czopska@cms-cmno.com

Wiktoria Murphy
Lawyer
T +48 22 520 82 68
M wiktoria.murphy@cms-cmno.com

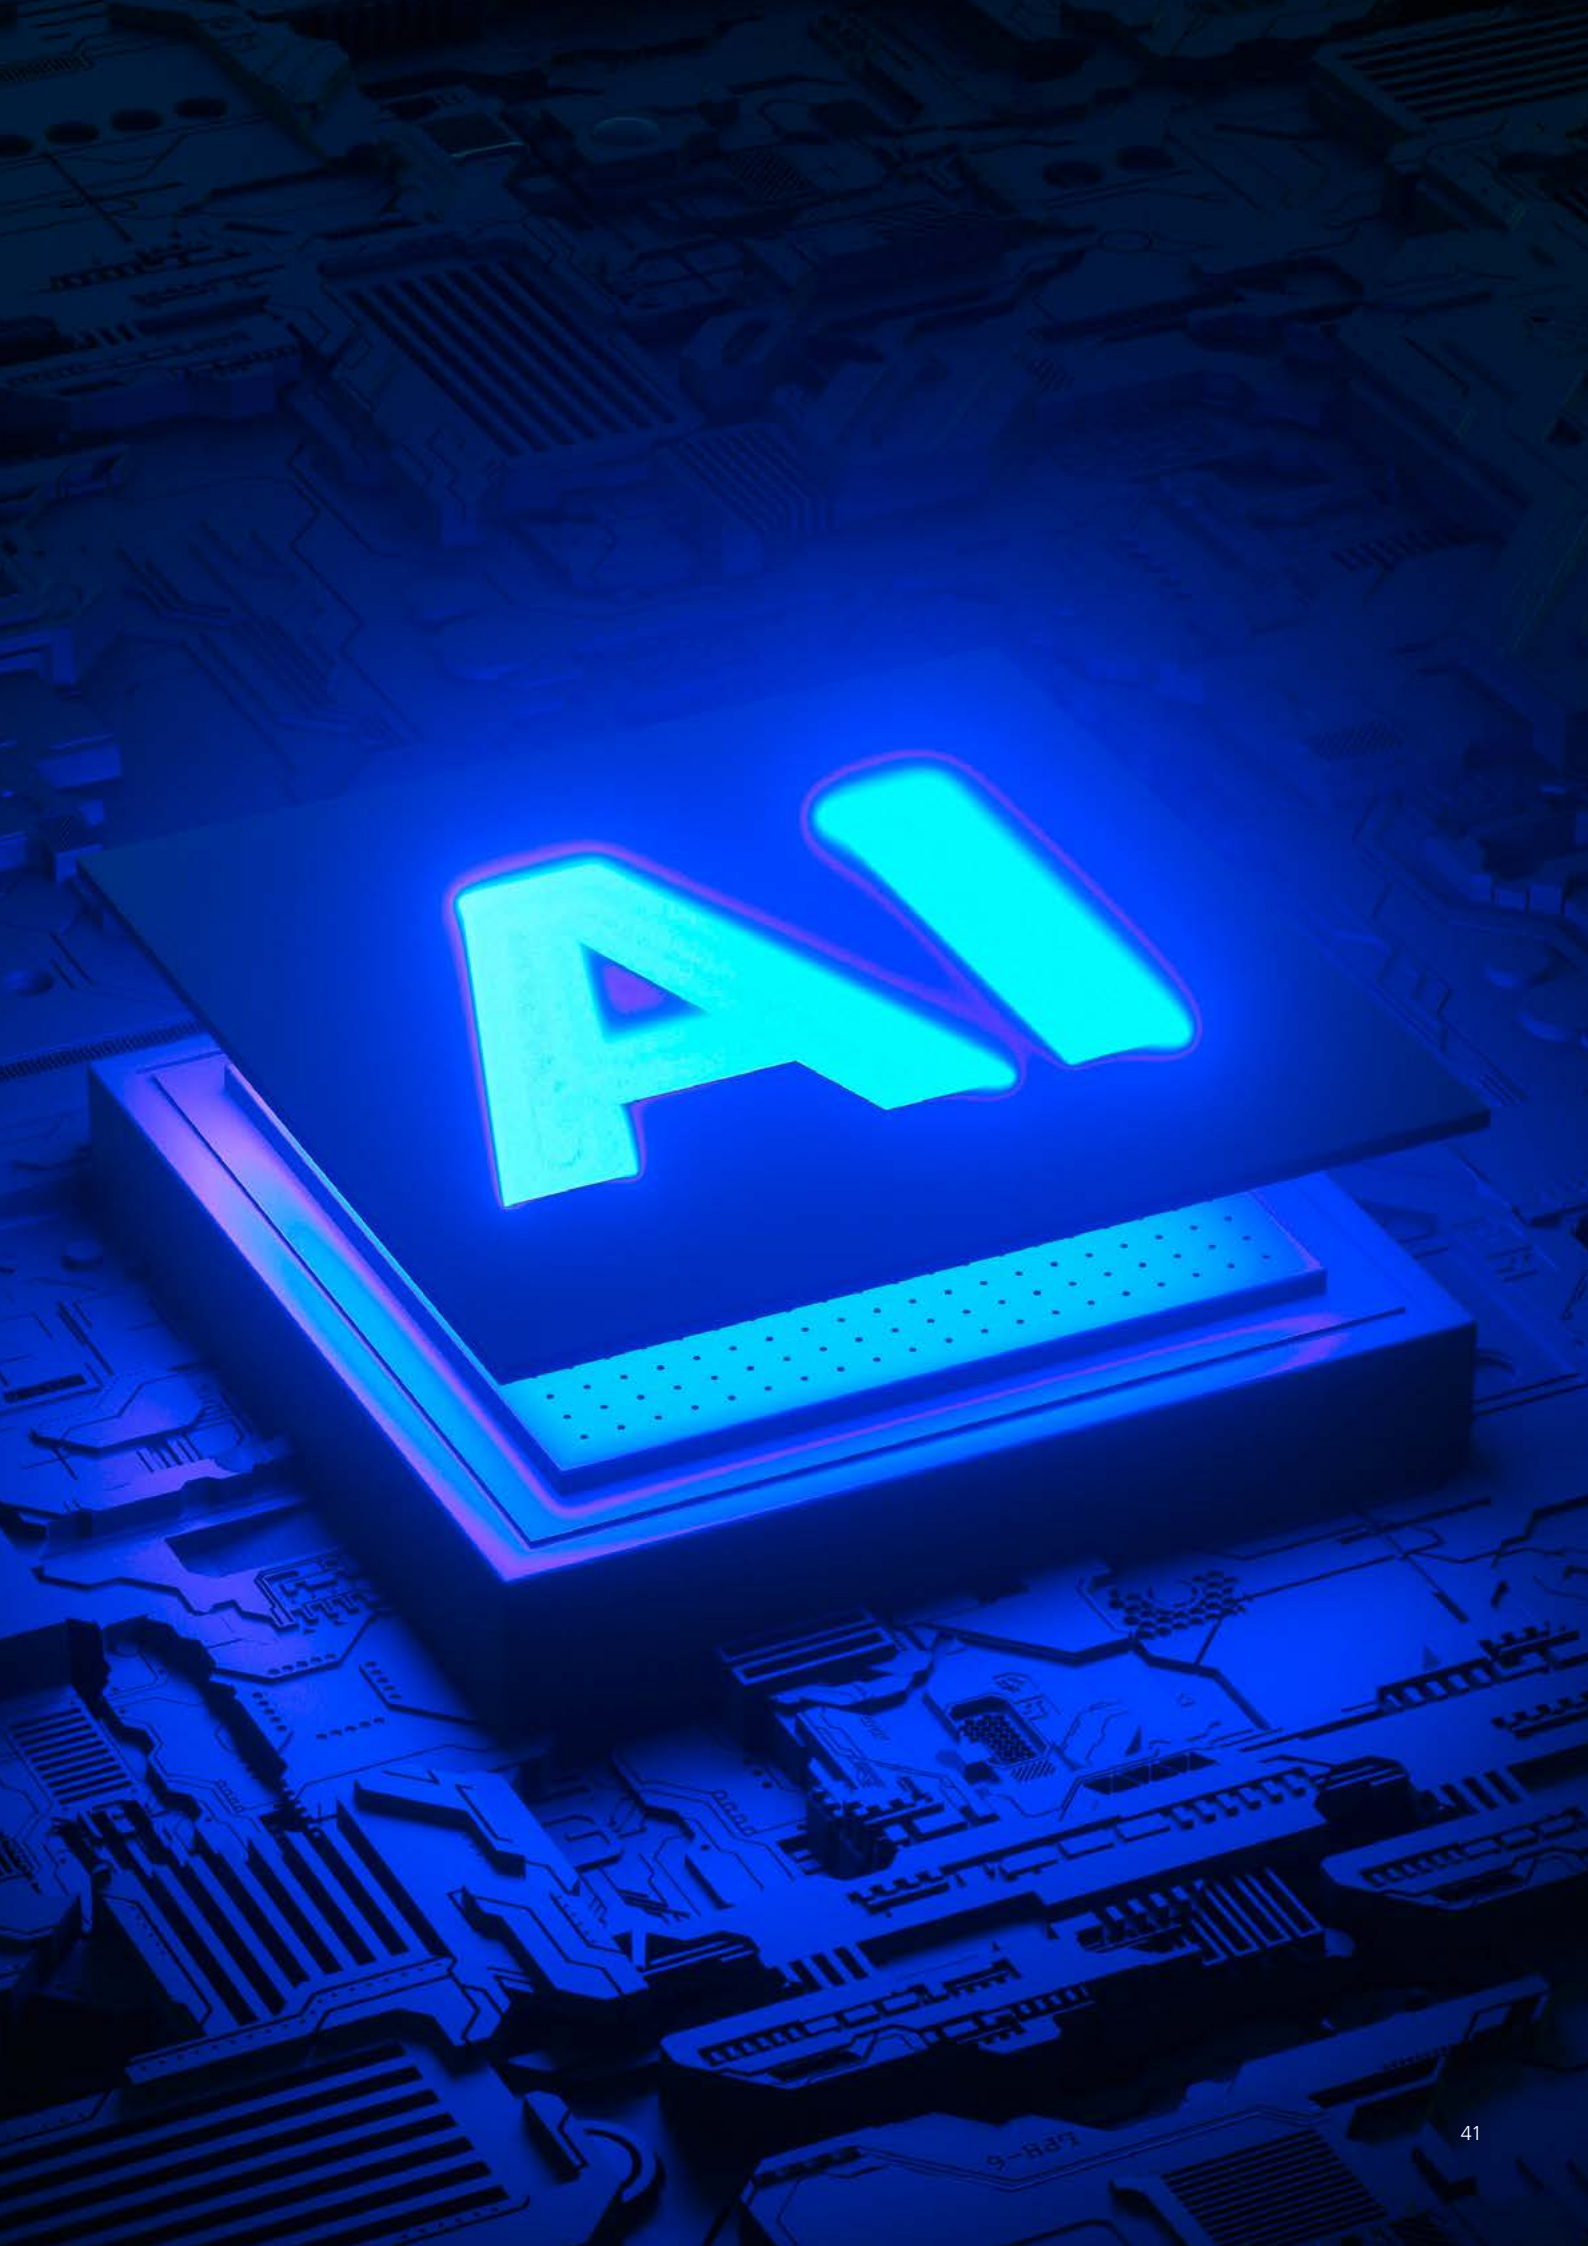
AI

CMS Legal Updates subscription service

Sign up now for the free online email alert service delivering commentary, analyses and insights from CMS experts on the legal issues affecting your business, directly to your inbox.
cms.law/subscription

The information held in this publication is for general purposes and guidance only and does not purport to constitute legal or professional advice.

CMS is an international organisation of independent law firms ("CMS Member Firms"). CMS LTF Limited (CMS LTF) is a company limited by guarantee incorporated in England & Wales (no. 15367752) whose registered office is at Cannon Place, 78 Cannon Street, London EC4N 6AF United Kingdom. CMS LTF coordinates the CMS Member Firms. CMS LTF provides no client services. Such services are solely provided by CMS LTF's CMS Member Firms in their respective jurisdictions. CMS LTF and each of its CMS Member Firms are separate and legally distinct entities, and no such entity has any authority to bind any other. CMS LTF and each CMS Member Firm are liable only for their own acts or omissions and not those of each other. The brand name "CMS" and the term "firm" are used to refer to some or all of the CMS Member Firms or their offices; details can be found under "legal information" in the footer of cms.law.

CMS locations:
Aberdeen, Abu Dhabi, Amsterdam, Antwerp, Barcelona, Beijing, Belgrade, Bengaluru, Bergen, Berlin, Bogotá, Bratislava, Brisbane, Bristol, Brussels, Bucharest, Budapest, Casablanca, Chennai, Cologne, Dubai, Dublin, Duesseldorf, Ebene, Edinburgh, Frankfurt, Funchal, Geneva, Glasgow, Gothenburg, Gurugram, Hamburg, Hong Kong, Hyderabad, Istanbul, Johannesburg, Kyiv, Leipzig, Lima, Lisbon, Liverpool, Ljubljana, London, Luanda, Luxembourg, Lyon, Madrid, Manchester, Maputo, Mexico City, Milan, Mombasa, Monaco, Mumbai, Munich, Muscat, Nairobi, New Delhi, Oslo, Paris, Podgorica, Poznan, Prague, Reading, Rio de Janeiro, Riyadh, Rome, Santiago de Chile, São Paulo, Sarajevo, Shanghai, Sheffield, Silicon Valley, Singapore, Skopje, Sofia, Stavanger, Stockholm, Strasbourg, Stuttgart, Sydney, Tel Aviv, Tirana, Vienna, Warsaw, Zagreb and Zurich.

Further information can be found at **cms.law**

© CMS Legal Services EEIG (2026) – 2606-0208939-9