




IRELAND

AI Laws of the World

Introduction



Artificial intelligence (AI) is rapidly transforming the way we work. As AI technology advances, governments are grappling with how to regulate its development to maximise benefits whilst mitigating risks.

Some jurisdictions have implemented AI-specific laws, whilst others are in the process of drafting regulations, and many rely on existing legal frameworks to address AI-related concerns.

This guide provides an overview of AI laws and proposed regulations across 50+ territories. It highlights key legislative developments, including regulations, proposed bills and guidelines issued by governmental bodies.


As the global landscape of AI regulation evolves at an accelerating pace, our 2026 Q3 snapshot provides an insightful overview of these developments, as well as of the common thematic approaches of lawmakers and AI-focused organisations around the world.

Global key contacts



Gareth Stokes

Partner
Global Co-Chair, AI Practice
gareth.stokes@dlapiper.com
[Full bio](#)



Danny Tobey

Partner
Global Co-Chair AI and
Data Analytics
danny.tobey@us.dlapiper.com
[Full bio](#)



Law / proposed law

Regulation (EU) 2024/1689 of the European Parliament and of the Council on harmonised rules on artificial intelligence (the **EU AI Act**) was published in the Official Journal of the European Union on 12 July 2024. It entered into force on 1 August 2024, although many of its provisions come into force on specific dates:

- 2 February 2025: General provisions and provisions relating to prohibited AI practices and AI literacy (Chapters I and II).
- 2 August 2025: Provisions relating to general-purpose AI (**GPAI**) models, governance and competent authorities (Chapters V and VII), and penalties (Chapter XII).
- 2 August 2026: Original enforcement date for most of the AI Act's remaining provisions, including the requirements for Annex III high-risk AI systems, AI regulatory sandboxes (Article 57(1)), and certain transparency obligations relating to chatbots, deepfakes and AI-generated content. However, under the Digital Omnibus for AI (defined below), the timeline for Annex III high-risk AI systems is extended to 2 December 2027, with a further extension to 2 August 2030 for high-risk AI systems intended to be used by public authorities. Similarly, the watermarking obligations imposed on generative AI content (Article 50(2)) are extended to 2 December 2026.
- 2 August 2027: Original enforcement date for provisions relating to high-risk AI systems that are safety components of products or products themselves (i.e., AI systems covered by Annex I). However, under the Digital Omnibus for AI (defined below), the timeline for Annex I high-risk AI systems is extended to 2 August 2028, with a further extension to 2 August 2030 for high-risk AI systems intended to be used by public authorities.

A new EU Product Liability Directive, Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products (**Product Liability Directive**), was published in the Official Journal of the European Union on 18 November 2024 and entered into force on 8 December 2024. Member States have until 9 December 2026 to implement the Product Liability Directive into national law. The Product Liability Directive modernises the EU-level strict product

liability regime, preserving the core principles of the previous law while adapting to new technologies by extending the scope to include software and AI. This regime is still limited to certain types of damages and applies only to consumers and other natural persons.

On 24 July 2026, the Council of the European Union adopted the proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as regards the simplification of the implementation of harmonised rules on AI (**Digital Omnibus for AI**). The Digital Omnibus for AI was published in the Official Journal of the European Union on 24 July 2026 (OJ L, 2026/1744) and entered into force on 27 July 2026. Key elements include:

- a new prohibition on non-consensual intimate AI-generated content and child sexual abuse material;
- extension of small and medium-sized enterprise (SME) regulatory exemptions to small mid-caps;
- clarification of the AI Office's supervisory powers over general-purpose AI models, including enforcement powers over AI systems embedded in very large online platforms (VLOPs) and very large online search engines (VLOSEs), and centralisation of enforcement of certain GPAI systems within the AI Office; national competent authorities retain jurisdiction over law enforcement, border management, the judiciary, and financial institutions;
- clarification of the AI literacy obligation;
- reduced administrative burden (e.g., registration simplification for non-high-risk systems falling into Article 6(3) exemptions); and
- extension of access to AI regulatory sandboxes, including through an EU-level sandbox, with national competent authorities required to establish AI regulatory sandboxes by 2 August 2027.

The Commission's 2022 proposal on adapting non-contractual civil liability rules to artificial intelligence (COM(2022) 496) (**AI Liability Directive**) was withdrawn under the Commission Work Programme 2025 (COM(2025) 45 final, Annex IV), with the withdrawal formalised in the Official Journal in October 2025. No replacement proposal is currently tabled. As proposed, the AI Liability Directive would have complemented the EU AI Act and the revised Product Liability Directive by harmonising national civil liability regimes and rules on the burden of proof in relation to AI. However, this ambitious project was ultimately abandoned, notably due to the significant divergence across Member States' legal frameworks.

On 22 June 2026, the European Parliament and Council reached a provisional agreement on a Directive aimed at combating the sexual abuse and sexual exploitation of children and child sexual abuse material (**CSAM**) (**CSAM Directive**). The CSAM Directive updates existing offences and introduces new criminal offences to address technological developments and emerging forms of CSAM, including the criminalisation of the production, possession, and dissemination of instructions on how to commit child sexual abuse or produce CSAM. The CSAM Directive also establishes offences related to AI, making it unlawful to design, adapt, acquire, possess, or distribute AI systems intended to generate CSAM. Additional offences include paying for access to livestreamed child sexual abuse, livestreaming such

abuse, organising travel for the purpose of child sexual abuse, and sexual extortion of children through threats to disclose CSAM. The Directive expands provisions on grooming, criminalising soliciting a child to produce or share CSAM, and increases penalties for offences such as the acquisition, possession, or access to CSAM, as well as its distribution or dissemination.

On 3 June 2026, the European Commission published a proposal for a Regulation establishing a framework of measures for strengthening Europe's cloud and AI ecosystem (**Cloud and AI Development Act** or **CADA**) (COM(2026) 502 final). CADA is addressed at the EU's limited computing capacity and dependence on non-European cloud providers, and aims to increase sustainable computing capacity, ensure data sovereignty and operational continuity, and enhance public sector resilience. While the EU AI Act regulates AI systems and GPAI models, CADA focuses on cloud and AI computing services and their underlying infrastructure. It establishes a four-level Union cloud computing sovereignty framework, with higher levels requiring that data generated by audited services cannot be used to train AI systems operated by entities outside the European Economic Area (EEA). The proposal also establishes 'Cloud and AI Leadership Initiatives' supporting frontier, physical, and industrial AI development. The AI Board would coordinate AI adoption activities under CADA.

AI compliance in Ireland

Ireland's AI regulatory landscape has entered a new phase with the enactment of the **Regulation of Artificial Intelligence Act 2026 (No 31 of 2026) (Irish AI Act)**, which was signed into law on 21 July 2026, and which came into operation on 31 July 2026 (other than paragraphs (a) to (e) of the definition of "applicable market surveillance authority" in section 78).

The Irish AI Act establishes the Oifig IS na hÉireann (AI Office of Ireland) (**Office**) as an independent statutory entity that will act as the central coordinating body for the implementation and enforcement of the EU AI Act in Ireland. The functions of the Office include, amongst others:

- coordination of the activities of relevant competent authorities to encourage consistent implementation and supervision of the Irish AI Act;
- facilitating cooperation and information sharing between relevant competent authorities in Ireland and supporting joint investigation and enforcement; and
- promoting and fostering AI innovation and literacy.

The Office is also empowered to coordinate with other similar international bodies, is responsible for maintaining a national AI register and may establish AI regulatory sandboxes.

In Ireland, the Product Liability Directive has not yet been implemented into law, but it is expected to be implemented ahead of the 9 December 2026 deadline.

Regulatory guidance / voluntary codes

In order to ensure the consistent, effective, and uniform application of the EU AI Act across the European Union, the European Commission has published various guidelines and codes of practice. These are non-binding, as only the Court of Justice of the European Union has authoritative interpretation powers. Codes of practice are voluntary compliance tools which help demonstrate compliance with binding AI Act obligations.

The Commission has published the following guidelines:

- 4 February 2025: [Guidelines on prohibited AI practices](#);
- 6 February 2025: [Guidelines on AI system definition](#);
- 18 July 2025: [Guidelines on the scope of obligations for general-purpose AI model providers](#);
- 19 May 2026: [Draft guidelines on the classification of high-risk AI systems \(High-Risk Classification Guidelines\)](#) (see the High-risk AI section);
- 20 July 2026: [Guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of the EU AI Act \(Article 50 Guidelines\)](#) (see the User transparency section).

The Commission has also published the following codes of practice:

- 10 July 2025: [General-purpose AI Code of Practice](#); and
- 10 June 2026: [Code of Practice on Transparency of AI-Generated Content \(Transparency Code of Practice\)](#) (see the User transparency section).

Under Article 95 of the EU AI Act, providers of AI systems that do not fall under the high-risk classification, as well as deployers, have the possibility to adopt voluntary codes of conduct in order to adopt, on a non-binding basis, technical solutions and industry best practices. Because of this, it is expected that the AI Office will issue further codes of conduct for this purpose.

To provide organisations with support identifying and implementing AI literacy initiatives, on 4 February 2025 the Commission launched a [repository of AI literacy practices](#).

In May 2024, the Council of Europe adopted the [Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law \(Framework Convention\)](#) (CETS No. 225). The Framework Convention was opened for signature on 5 September 2024 in Vilnius, and the European Union signed the same day pursuant to Council Decision (EU) 2024/2218. It is an international, legally binding treaty aiming to ensure that activities within the lifecycle of AI systems are fully consistent with human rights, democracy, and the rule of law, whilst being conducive to technological progress and innovation.

The Irish government and certain competent authorities in Ireland have also published guidance on the use and adoption of AI, including those detailed below.

Guidelines for the Responsible Use of AI in the Public Service (2025)

On 7 May 2025, the Irish Government (specifically, the Department of Public Expenditure, Infrastructure, Public Service Reform and Digitisation) published *Guidelines for the Responsible Use of AI in the Public Sector (Public Sector Guidelines)*, which were developed to empower public servants to use AI in the delivery of services. The Public Sector Guidelines set the overarching framework for adoption and use of AI by the public service, and are built around 'Seven Principles for Responsible AI', namely:

- Human agency and oversight
- Technical robustness and safety
- Privacy and data governance
- Transparency
- Diversity, non-discrimination and fairness
- Societal and environmental well-being
- Accountability

National Cyber Security Centre (NCSC) Guidance on AI

To complement the Public Sector Guidelines, the National Cyber Security Centre (NCSC) in Ireland has published cyber security guidelines for AI deployments (*Cyber Security Guidelines*), which address the cyber security obligations arising from the use and adoption of AI.

The Cyber Security Guidelines are directed at public bodies, but the NCSC recognises within these guidelines that they may also be useful for management boards who will have expanded responsibilities under the Network and Information Security Directive 2 (NIS2). They are structured around seven principles spanning the phases of the AI lifecycle as follows:

- Design – Principle 1: Build trust and security through informed design and risk awareness.
- Development – Principle 2: Identify and secure all assets ensuring end-to-end traceability and recovery. Principle 3: Build resilience through effective security controls and supply chain assurance.
- Deployment – Principle 4: Implement rigorous testing and validation to ensure reliability. Principle 5: Support oversight and assurance through accountable, explainable and transparent operation.
- Maintenance – Principle 6: Maintain secure operations through continuous oversight and rapid response.

- End-of-life – Principle 7: Safely retire artefacts and document that compliance obligations have been met.

Central Bank of Ireland Regulatory and Supervisory Outlook Report 2026

Pursuant to European Union (Artificial Intelligence) (Designation) Regulations 2025, the Central Bank of Ireland (CBI) is designated as the market surveillance authority in Ireland, responsible for the financial supervision of institutions placing on the market, putting into service, or using AI systems in direct connection with the provision of financial services.

In February 2026, the CBI published its regulatory and supervisory outlook report (**Report**) in which it provides guidance for the use of AI in financial services. The Report states that firms must adhere to the following standards when deploying AI:

- Strategic alignment: Firms must ensure their use of AI is appropriate for the specific business challenge being addressed.
- Accountability and explainability: They should establish clear accountabilities and responsibility, human oversight of decisions and their explainability.
- Proportionate governance: Risk management practices must be commensurate with the scale, scope and sensitivity of the AI deployment.
- Compliance: Processes should be in place to ensure all EU AI Act obligations are met, including transparency requirements.

Appointed supervisory authority

European Level

The European Commission established the European AI Office (AI Office) on 24 January 2024. The AI Office is a European Commission function and forms part of the Directorate-General for Communications Networks, Content and Technology; it must therefore operate in accordance with the Commission's internal processes. The AI Office is responsible for assisting the European Commission with the oversight, monitoring, and enforcement of requirements for GPAI models and systems. Under the Digital Omnibus for AI, the AI Office has been granted enforcement powers over AI systems embedded in VLOPs and VLOSEs, and certain GPAI systems enforcement is centralised within the AI Office. It is primarily made up of hired full-time staff from a range of backgrounds such as technology specialists, economists, policy specialists, and lawyers.

In addition, the European Artificial Intelligence Board (AI Board) has been established under Article 65 of the EU AI Act. The AI Board's core responsibility is to advise and assist the Commission and Member States to facilitate the consistent and effective application of the EU AI Act. The AI Board includes a representative from each Member State, and the AI Office and the European Data Protection Supervisor participate as non-voting observers.

Member State Level

Article 70 of the EU AI Act concerns the designation of national competent authorities by EU Member States. It specifies that each Member State shall establish or designate as national competent authorities at least one notifying authority and at least one market surveillance authority for the purposes of the general supervision and enforcement of the EU AI Act. Where multiple market surveillance authorities are appointed, one of the market surveillance authorities must act as the single point of contact. The authorities must operate independently and without bias. Member States were required to notify the Commission of their appointed authorities and publicly available information on how to contact them by 2 August 2025, and must report to the Commission on the financial and human resources of their competent authorities by the same date and every two years thereafter (Article 70(6)). The designation of national authorities has been slower than expected, which may affect early enforcement consistency.

Supervisory authority in Ireland

The Irish AI Act establishes the Oifig IS na hÉireann (AI Office of Ireland) (**Office**) as an independent statutory entity that will act as the central coordinating body for the implementation and enforcement of the EU AI Act in Ireland, and will be the single point of contact for the purposes of Article 70(2) of the EU AI Act.

The Office is responsible for the coordination of the activities of relevant competent authorities to: (i) encourage consistent implementation and supervision of the Irish AI Act; (ii) facilitate cooperation and information sharing between relevant competent authorities in Ireland; and (iii) support joint investigation and enforcement by relevant competent authorities.

For the purposes of Article 70(1) of the EU AI Act, Ireland has designated:

- The Minister for Enterprise, Tourism and Employment as a national competent authority; and
- Each of the following market surveillance authorities as national competent authorities for specified purposes:
 - Health and Safety Authority
 - Competition and Consumer Protection Commission
 - Marine Survey Office of the Department of Transport
 - Commission for Communications Regulation
 - Commission for Railway Regulation
 - Health Products Regulatory Authority

The Central Bank of Ireland is designated as the market surveillance authority for the purposes of Article 74(6) of the EU AI Act.

The Data Protection Commission is designated as the market surveillance authority for the purposes of Article 74(8) of the EU AI Act.

The Irish AI Act also appoints additional market surveillance authorities, as follows:

- (a) Coimisiún na Meán (appointment not yet in force);
- (b) the Commission for Railway Regulation;
- (c) the Commission for Communications Regulation;
- (d) the Commission for Regulation of Utilities;
- (e) the Data Protection Commission (appointment not yet in force);
- (f) the Health and Safety Authority;
- (g) the Health Products Regulatory Authority;
- (h) the Health Service Executive;
- (i) the Marine Survey Office of the Department of Transport;
- (j) the National Transport Authority; and
- (k) the Workplace Relations Commission.

For the purposes of Article 28 of the EU AI Act, Ireland has designated four notifying authorities, each in respect of specified points:

- The Minister for Enterprise, Tourism and Employment
- The Minister for Transport
- The Health Products Regulatory Authority
- The Commission for Communications Regulation

Pursuant to Article 77, Ireland has also designated nine national public authorities to supervise and enforce in respect of fundamental rights. Their role is to exercise their existing mandates in circumstances involving the use of AI systems, with enhanced rights to facilitate that work. These are:

- The Data Protection Commission
- Coimisiún na Meán (Media Commission)
- Irish Human Rights and Equality Commission
- An Coimisiún Toghcháin (Electoral Commission)
- Ombudsman for Children's Office
- Ombudsman for the Defence Forces
- Financial Services & Pensions Ombudsman
- Environmental Protection Agency Ombudsman

Definitions

AI System

Article 3(1) of the EU AI Act defines an 'AI system' as follows:

"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 input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments".

The EU AI Act uses a technology-neutral definition, focusing on the effect of the system rather than the techniques used. There are several key features of the definition which, acting together, distinguish the AI system from more traditional software systems. The central characteristics are the level of autonomy and adaptiveness in how the system operates and the ability for the system to infer how to generate outputs. So, an AI system must be able to operate independently at some level (like many existing technologies) but must also be able to apply logic to draw conclusions from data it is given. It may also adapt after deployment, in effect by continuing to 'learn'. These features are more akin to human capability than traditional technology systems, which operate using more fixed and pre-determined paths to process data. These outputs must influence physical or virtual environments, whether by making decisions or through other means.

The EU AI Act also sets out specific rules for GPAI models. GPAI models differ from AI systems; they can be an essential component integrated into an AI system, but do not themselves constitute an AI system until further components are added (such as an interface). For more information, see the [Controls on generative AI section](#).

Provider

Article 3(3) of the EU AI Act defines a 'provider' as follows:

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

Those falling within this definition as a 'provider' have significant responsibility for ensuring compliance with the EU AI Act, and so identifying the provider will be crucial for businesses and may well influence their choice of business/deployment model.

The provider is responsible for putting the AI on the market either by making it first available in the market or by directly putting the AI into use for its own purposes and under its own name or trademark. An organisation may also become a downstream provider if it makes substantial modifications to a system or changes its intended purpose (Article 25(1)). Guidance from the European Commission is expected on what counts as a 'substantial modification'. At this stage, the only conclusive criterion is that such modification must not have been foreseen by the provider in the initial conformity assessment carried out by the provider.

Payment is not relevant, which will impact GPAI models supplied onto the market on an open-source basis or under free commercial terms.

Deployer

Article 3(4) of the EU AI Act defines a 'deployer' as follows:

"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 simple terms, a 'deployer' is an entity that uses an AI system other than for personal, non-professional use. Although the burden of responsibility on a deployer is not as great as on providers, there are still obligations that it must fulfil.

Note that the EU AI Act also implements requirements for organisations performing other roles (as distributor, importer, product manufacturer, and authorised representative). Together with the deployer and provider, such organisations are referred to as 'operators' of AI. Importantly, the same operator may qualify simultaneously as more than one of these roles if they meet the respective conditions. For instance, it is possible to be both the provider and the deployer of an AI system at the same time.

Prohibited activities

Certain AI practices are banned outright under Article 5 of the EU AI Act due to their potential for harm and ethical concerns. These prohibitions aim to protect EU citizens from the most intrusive and potentially abusive uses of AI.

Under Article 5, these uses and technologies include:

- **Subliminal techniques:** Deploying subliminal techniques or techniques that are manipulative or deceptive and have the effect or objective of materially distorting a person's behaviour by impairing their ability to make an informed decision, causing

them to make a decision they would not otherwise have taken, in a manner that causes significant harm to them or others (or is reasonably likely to do so).

- **Exploiting vulnerabilities:** Exploiting vulnerabilities of specific groups due to age, disability, or social or economic situation – as with subliminal techniques, this must have the effect or objective of materially distorting behaviour and cause significant harm to them or others (or be reasonably likely to do so).
- **Social scoring:** Evaluating or classifying natural persons or groups based on their social behaviours or personality characteristics (known, inferred, or predicted) leading to either or both, unfavourable treatment of them or others in social contexts unrelated to the context in which the data was originally gathered or that is unjustified or disproportionate to their social behaviour or its gravity.
- **Crime profiling:** Assessing the risk of an individual committing a crime, based on the profiling of that person and assessing their personality traits (as opposed to using such systems to support a human assessment of the involvement of a person).
- **Facial recognition databases:** Creating or expanding facial recognition databases through untargeted scraping of facial images from the internet or CCTV footage.
- **Inferring emotions:** Inferring emotions in workplaces and educational institutions (except for medical or safety reasons).
- **Biometric categorisation:** Categorising natural persons based on their biometric data to deduce or infer sensitive information about them (i.e., their race, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation) except where based on lawfully acquired datasets (including in law enforcement).
- **Biometric identification:** Engaging in 'real-time' biometric identification systems in publicly accessible spaces for law enforcement purposes, except under specific exempt circumstances.
- As a new prohibited practice introduced by the Digital Omnibus for AI, **Non-consensual intimate AI-generated content and CSAM:** Placing on the market, putting into service, or using AI systems that (i) are intended to create, generate, or manipulate, or (ii) are designed or capable such that the creation, generation, or manipulation is a reasonably foreseeable outcome without adequate technical safeguards, non-consensual intimate content of natural persons (including non-consensual intimate deepfakes) or child sexual abuse material. Companies must comply with this prohibition by 2 December 2026.

Prohibited activities in Ireland

Recital 40 of the EU AI Act sets out Ireland's exemptions from certain prohibitions on AI practices in the context of police cooperation and judicial cooperation in criminal matters. Such exemptions arise from historic features of EU treaty law, namely Article 6a of Protocol No 21 and cover:

- Article 5(1)(g) – The use of AI systems to deduce sensitive information about individuals;

- Article 5(1)(d) – The use of AI systems to predict the risk of an individual committing a criminal offence; and
- Article 5(1)(h), Article 5(2) to 5(6) and Article 26(10) – The use of 'real-time' remote biometric identification systems in public spaces for law enforcement.

The Irish Data Protection Commission (DPC) has been active in opening a number of consultations and investigations on certain AI systems.

- In 2026, the DPC opened an inquiry into X Internet Unlimited Company under Section 110 of the Data Protection Act 2018 regarding the use of personal data for AI training which was not in compliance with Articles 5, 6, 25 and 35 General Data Protection Regulation (GDPR) obligations.
- In 2025, the DPC engaged with LinkedIn regarding the training of a proprietary generative AI model using the personal data of users. As a result of this consultation, LinkedIn amended its plans.
- In 2024, the DPC ordered X to suspend training of an AI chatbot after issuing High Court proceedings pursuant to Section 134 of the Data Protection Act 2018.
- In 2024, the DPC launched a statutory inquiry into Google Ireland under Section 110 of the Data Protection Act 2018, regarding compliance with data protection obligations pursuant to Article 35 of the GDPR for its Pathways Language Model 2 (PaLM 2).
- In 2023/24, the DPC engaged with Meta in relation to the training of its large language model (LLM) using public content shared across the EU, which led to the DPC seeking a formal GDPR opinion on the matter from the European Data Protection Board (EDPB). Engagement with Meta and oversight of its implemented measures and improvements is ongoing.

In February 2025, the DPC also became one of five signatory data protection authorities on the Paris declaration to reaffirm their commitment to implementing data governance that promotes innovative and privacy-protecting AI. At the 47th Global Privacy Assembly in Seoul in September 2025, this declaration was subsequently expanded, and the number of signatories grew from five to 20.

High-risk AI

Article 6 of the EU AI Act sets out classification rules for high-risk AI systems, stating that high-risk AI systems fall within two categories: (i) safety components of products or products themselves regulated by existing EU product safety laws (listed in Annex I, e.g., medical devices, automotive AI); or (ii) systems used in specified areas (listed in Annex III), namely:

- **Critical infrastructure:** AI systems used as safety components in the management or operation of critical digital infrastructure, road traffic, or the supply of water, gas, heating, or electricity.

- **Education and vocational training:** AI systems that determine access to education or training or otherwise impact a person's future opportunities and career development, and AI systems used for monitoring and detecting prohibited behaviour during tests.
- **Employment and worker management:** AI systems used in hiring (including the placement of targeted job advertisements), performance evaluation, promotion, or termination decisions.
- **Access to essential private and public services:** AI systems that evaluate eligibility for essential public services, such as social security and healthcare, as well as AI systems for evaluating and classifying emergency calls and dispatching emergency services. Additionally, AI systems used to evaluate creditworthiness or during the risk assessment and pricing of life and health insurance.
- **Law enforcement:** AI systems used by law enforcement for risk assessments, predicting criminal activities (the risk of individuals becoming victims of crime, risk of (re-)offending, or otherwise during criminal investigations), for polygraphs (i.e., 'lie detectors' or similar tools), and assessing reliability of evidence.
- **Border control and migration:** AI systems used to assess visa applications, asylum claims, and border security, including for polygraphs (i.e., 'lie detectors' or similar tools), and for detecting, recognising, or identifying individuals in migration contexts.
- **Judicial and democratic processes:** AI systems assisting judicial authorities with researching and interpreting facts and the law and applying the law to a set of facts, as well as AI systems used for influencing the outcome of elections or referendums or voting behaviour.
- **Biometric identification and categorisation:** AI systems that perform remote biometric identification, are used to categorise individuals based on biometric data or other sensitive or protected attributes, or are used for emotion recognition purposes.

Chapter III, Section 2 of the EU AI Act (Articles 9 to 15) imposes mandatory requirements on high-risk AI systems, including risk management systems (Article 9), data governance (Article 10), technical documentation (Article 11), record-keeping and logging (Article 12), transparency and information to deployers (Article 13), human oversight measures (Article 14), and accuracy, robustness, and cybersecurity (Article 15). Providers must also implement quality management systems (Article 17) and post-market monitoring (Article 72).

The European Commission has the power to amend the abovementioned categories of high-risk AI systems, including to modify any existing use cases or add new ones (Article 7(1) of the EU AI Act).

The Digital Omnibus for AI clarifies the definition of 'safety component' (Article 3(14)), specifying that an AI system fulfils a safety function only where its intended purpose is to prevent or mitigate risks to health and safety. AI systems used solely for non-safety-related functions such as user assistance, performance optimisation, service efficiency, automation, convenience, or quality control do not qualify as safety components.

The 19 May 2026 High-Risk Classification Guidelines are intended to assist providers, deployers, and market surveillance authorities in determining whether an AI system

qualifies as high-risk. They provide interpretation of key classification concepts and practical examples.

Where an AI system falls into one of the two abovementioned categories of high-risk AI systems but does not pose significant risk of harm to health, safety, or fundamental rights, the operators of such AI systems are relieved from the requirements imposed for high-risk AI systems (except for the EU database registration). However, to benefit from such exemption, a thorough assessment must be documented and strict conditions must be met (the 19 May 2026 High-Risk Classification Guidelines published in May 2026 provide interpretive guidance). Importantly, providers claiming this exemption under Article 6(3) must still register the AI system in the EU database pursuant to Article 49(2). The Digital Omnibus for AI maintains this mandatory registration requirement for effective market surveillance and public accountability, but simplifies the registration process.

The Digital Omnibus for AI extends the deadlines for high-risk AI system requirements. Specifically:

- for Annex III high-risk AI systems, the requirements apply from 2 December 2027 (extended from 2 August 2026);
- for Annex I high-risk AI systems, the requirements apply from 2 August 2028 (extended from 2 August 2027); and
- for high-risk AI systems intended to be used by public authorities, providers and deployers must comply with the EU AI Act requirements by 2 August 2030.

Controls on generative AI

General-Purpose AI Models

Article 3(63) of the EU AI Act defines a GPAI (general-purpose AI) model as 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, except AI models that are used for research, development or prototyping activities before they are placed on the market."

GPAI models are versatile and can be applied across various domains and contexts. The Act sets requirements to ensure that these specific models, due to their broad applicability and the wide range of tasks they can complete, adhere to high ethical and safety standards. Please note that not all AI models are GPAI models, and the EU AI Act only regulates the latter.

General-Purpose AI Models with Systemic Risk

Article 3(65) of the EU AI Act defines 'systemic risk' as:

"a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain".

Article 51 of the EU AI Act classifies a GPAI model as having systemic risk if it has high-impact capabilities (this is presumed when the cumulative amount of computation used for training exceeds 10^{25} floating point operations (FLOPs), though the Commission may also designate models based on other indicators and benchmarks) or based on a decision of the Commission.

Systemic risk involves the broader, cumulative impact of GPAI models on society. This encompasses scenarios where GPAI models could lead to significant disruptions or risks, necessitating a regulatory focus to prevent widespread adverse effects and ensure resilience across sectors. In view of the higher risks, the Act sets additional requirements for GPAI models with systemic risk.

Importantly, the requirements of a GPAI model/system (i.e., without a specific use case) and the requirement of an AI system based on its risk profile (depending on the use case at stake) can be cumulative. For instance, if the provider of a GPAI model integrates its model in a high-risk AI system, then the rules for both GPAI models and high-risk AI systems should be complied with.

Enforcement / fines

The EU AI Act enforces compliance through a structured framework of fines and sanctions, varying in severity based on the nature of the non-compliance. Member States were required to lay down rules on penalties (including administrative fines) and notify them to the Commission by 2 August 2025.

For non-compliance with prohibited AI practices (see the Prohibited activities section), fines can reach up to EUR 35 million or 7% of the total worldwide annual turnover, whichever is higher.

Breaches of high-risk AI system requirements (see the High-risk AI section) can incur fines up to EUR 15 million or 3% of the total worldwide annual turnover.

Other non-compliance issues, such as providing incorrect or misleading information, can result in fines up to EUR 7.5 million or 1% of the total worldwide annual turnover. This applies to breaches not covered by the highest or significant sanctions. Reduced fine caps apply to SMEs and start-ups.

The Irish AI Act endorses the framework of fines and sanctions specified in the EU AI Act.

Pursuant to Article 99(8) of the EU AI Act, which provides for each Member State to lay down rules on to what extent administrative fines may be imposed on public authorities and bodies established in that Member State, the Irish AI Act provides that where an adjudicator finds that an administrative fine should be imposed on a public body (within the meaning of section 10 of the Data Sharing and Governance Act 2019), the amount of the fine is limited to EUR 1 million.

In Ireland, relevant market surveillance authorities may appoint persons to be authorised officers to assist in carrying out their functions under the Irish AI Act and the EU AI Act. Authorised officers are granted certain powers under the Irish AI Act pursuant to Article 70.

Chapter 2 of the Irish AI Act sets out the procedures authorised officers may take where they reasonably suspect that a person has infringed the Irish AI Act, including referral to an adjudicator. The AI Office of Ireland nominates adjudicators to be appointed by the Minister for Enterprise, Tourism and Employment. The Irish AI Act sets out the role and powers of an adjudicator following referral, the procedure to be followed and the administrative fines that may be imposed, which, as noted in the Enforcement/Fines section, aligns with the framework of fines and sanctions specified in the EU AI Act.

User transparency

Article 50 of the EU AI Act sets out transparency obligations for providers and deployers of certain AI systems, which include that:

- Article 50(1): providers of AI systems must ensure that natural persons using an AI system are informed that they are interacting with an AI system unless this is obvious to the natural person (this obligation excludes AI systems authorised by law to detect, prevent, investigate, or prosecute criminal offences).
- Article 50(2): providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video, or text content must ensure that the outputs are marked in a machine-readable format and detectable as artificially generated or manipulated. This obligation excludes AI systems that perform an assistive function for standard editing, AI systems that do not substantially alter the input data, and AI systems authorised by law to detect, prevent, investigate, or prosecute criminal offences. Providers must also process data in accordance with other relevant EU laws.

- Article 50(3): deployers of emotion recognition or biometric categorisation systems must inform the affected natural persons.
- Article 50(4): deployers of AI systems that generate or manipulate image, audio, or video content constituting deepfakes must disclose that the content has been artificially generated or manipulated. For content forming part of an evidently artistic, creative, satirical, fictional, or analogous work or programme, this disclosure must be made in an appropriate manner that does not hamper the display or enjoyment of the work.
- Article 50(5): deployers of AI systems that generate or manipulate text published to inform the public on matters of public interest must disclose that the text has been artificially generated or manipulated. This obligation does not apply where the AI-generated content has undergone human review or editorial control and a natural or legal person holds editorial responsibility for the publication.

These obligations apply from 2 August 2026 (except for the watermarking obligation under Article 50(2), for which providers of AI systems placed on the market before 2 August 2026 have until 2 December 2026 to comply).

The Transparency Code of Practice, published on 10 June 2026, is a voluntary instrument that sets out practical steps to help providers and deployers of generative AI systems meet the Article 50 transparency obligations. The Code is structured into two sections: Section 1 addresses provider obligations under Article 50(2), covering the marking and detection of AI-generated or manipulated content through machine-readable solutions; Section 2 addresses deployer obligations under Article 50(4) and (5), covering the labelling of deepfakes and AI-generated text published to inform the public on matters of public interest. On 8 July 2026, the Commission concluded that the Code adequately covers these obligations and facilitates their effective implementation.

The Article 50 Guidelines, published on 20 July 2026, provide practical guidance to assist providers, deployers, and competent authorities in ensuring compliance with the Article 50 transparency obligations. The Guidelines clarify key concepts, including what constitutes a directly interactive AI system (such as chatbots), the scope of synthetic content, and the treatment of deepfakes. They also address exceptions for AI systems that perform assistive functions for standard editing, and for AI systems authorised by law to detect, prevent, investigate, or prosecute criminal offences. The Guidelines complement the Transparency Code of Practice and explain how adherence to the Code may be used to demonstrate compliance with the EU AI Act's requirements.

Under the Digital Omnibus for AI, providers of generative AI systems that were already placed on the market before 2 August 2026 have a four-month transitional period to comply with the Article 50(2) watermarking obligations, ending on 2 December 2026.

Fairness / unlawful bias

At its core, the EU AI Act is driven by the imperative to safeguard the fundamental rights of EU citizens. The rapid advancement of AI technologies has introduced significant benefits but also potential risks, such as biases in decision-making systems and privacy infringements. The AI Act aims to mitigate these risks by establishing clear rules that ensure AI systems respect the rights enshrined in the EU Charter of

Fundamental Rights. This focus on human-centric AI seeks to enhance trust and acceptance among the public, thereby promoting wider adoption of AI technologies in a responsible manner.

Within the EU AI Act, non-discrimination and fairness are incorporated within the following:

- Recital 27 includes seven principles for trustworthy AI including ensuring that AI systems are developed and used in a way that includes diverse actors and promotes equal access, gender equality, and cultural diversity, while avoiding discriminatory impacts and unfair biases that are prohibited by Union or national law.
- Article 10 sets out data and data governance requirements for high-risk AI systems and includes a requirement to examine and assess possible bias in training, validation, and testing data sets.
- Deployers are required to ensure that any input data is relevant and sufficiently representative in view of the intended purpose of the high-risk AI system (Article 26 (4)).
- Under Article 4a (introduced by the Digital Omnibus for AI), providers and deployers may exceptionally process special categories of personal data (such as data revealing racial or ethnic origin, political opinions, or religious beliefs) where strictly necessary for bias detection and correction. This legal basis is subject to strict safeguards, including pseudonymisation, access controls, deletion once bias correction is complete, and documentation in processing records.
- Article 27 requires deployers of high-risk AI systems (where those deployers are public authorities or private entities providing public services) to conduct fundamental rights impact assessments before putting the system into use (see the High-risk AI section). These assessments must address, among other matters, the specific risks of harm to marginalised persons or vulnerable groups and the foreseeable impact on the categories of natural persons affected. The 19 May 2026 High-Risk Classification Guidelines provide further interpretive guidance on assessing when an AI system poses a significant risk to fundamental rights, including risks of discrimination, and clarify that AI systems used in contexts with heightened discrimination risks (such as employment, credit, and access to essential services) require careful assessment against the criteria in Article 6(3).
- Article 5 prohibits certain AI practices that are inherently discriminatory, including AI systems that categorise natural persons based on biometric data to deduce or infer sensitive information about them (such as race, political opinions, trade union membership, religious beliefs, sex life, or sexual orientation), except where based on lawfully acquired datasets (see the Prohibited activities section).

The Framework Convention addresses the issue of bias (most notably in Articles 10 and 11 relating to 'equality and non-discrimination') and highlights that AI has the potential to create and reinforce biases and that bias and discrimination by AI can cause manifest harm to individuals and to society. The Framework Convention encourages the development and sharing of strategies to counter these risks, including debiasing datasets in research and development and by the development of

rules on data processing. This approach has the potential to turn software, algorithms, and data into an asset in fighting bias and discrimination in certain situations, and a force for equal rights and positive social change.

Human oversight

Human oversight is crucial for preventing and mitigating risks associated with an AI system's operation. Providers must also ensure that operators are adequately trained to oversee the AI system, understand its functionalities, and respond appropriately to any issues. Effective human oversight enhances the safety and reliability of high-risk AI systems, ensuring they operate within acceptable parameters and can be controlled in case of unexpected behaviour or malfunctions.

Article 14 of the EU AI Act deals with human oversight, stating that providers must implement measures to ensure effective human oversight of high-risk AI systems. This involves designing the system with mechanisms that allow human operators to monitor, intervene, and deactivate the AI system if necessary. Providers of high-risk AI systems are required to ensure that systems falling under their responsibility are compliant with this requirement (Article 16(a)) and to include the human oversight measures within the 'instructions for use' for the high-risk AI system (Article 13(3)(d)).

In addition, deployers of high-risk AI systems are required to comply with the provider's 'instructions for use' and to assign human oversight to persons that have the necessary competence, training, and authority as well as necessary support (Article 26 (1) and (2)).

Finally, Recital 27 of the EU AI Act includes seven principles for trustworthy AI including ensuring that AI systems apply human agency and oversight. This means that AI systems are developed and used as a tool that serves people, respects human dignity and personal autonomy, and functions in a way that can be appropriately controlled and overseen by humans.

Key contacts



Mark Rasdale
Partner
DLA Piper
mark.rasdale@dlapiper.com
[View bio](#)



Claire O'Brien
Senior Associate
DLA Piper
claire.obrien@dlapiper.com
[View bio](#)



Maria Fitzpatrick

Legal Director

DLA Piper

maria.fitzpatrick@dlapiper.com

[View bio](#)



Conor McEneaney

Partner

DLA Piper

conor.mceneaney@dlapiper.com

[View bio](#)

For more information

To learn more about DLA Piper, visit dlapiper.com or contact:



Gareth Stokes

Partner

Global Co-Chair, AI Practice

gareth.stokes@dlapiper.com

[Full bio](#)



Danny Tobey

Partner

Global Co-Chair AI and Data Analytics

danny.tobey@us.dlapiper.com

[Full bio](#)

About us

DLA Piper is a global law firm with lawyers located in more than 40 countries throughout the Americas, Europe, the Middle East, Africa and Asia Pacific, positioning us to help companies with their legal needs around the world.

dlapiper.com