WilmerHale's Guide to the European Union's AI Act

WILMERHALE

A World First

The European Union's Artificial Intelligence Act (AI Act) is considered to be the world's first comprehensive horizontal legal framework for AI. It provides for EU-wide rules on data quality, transparency, human oversight, and accountability. With challenging requirements, significant extraterritorial effects, and fines of up to 35 million euros or 7% of global annual revenue (whichever is higher), the AI Act will have a profound impact on a significant number of companies conducting business in the European Union.

The Time to Prepare Is Now

The AI Act was published in the Official Journal of the European Union on July 12, 2024, as "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence." While the AI Act will generally apply starting on August 2, 2026, the exact milestones are quite nuanced and complex, with some provisions already applying since February 2, 2025.

Several categories of affected actors may face the need to significantly redesign their products and services, a process which should be initiated as soon as possible. Non-AI companies are subject to similar time constraints, as they will need to understand the technology and establish their own risk thresholds to effectively navigate compliance.

What You Will Find in This Guide

The European Union's AI Act is a long, complex, and technical text that is full of cross-references to other European legislative instruments and that uses concepts that sometimes require prior knowledge of European law, and of data protection law in particular.

This guide offers a simplified presentation of the AI Act's requirements, focusing on the most relevant aspects to help companies maintain compliance.

To this end, this guide covers the topics listed below.

1. Scope and Approach of the AI Act 6
2. Critical Milestones on the Road to Full Applicability of the AI Act 10
3. Prohibited AI Practices 12
4. Definition and Requirements for High-Risk AI Systems 15
5. Obligations for Deployers, Providers, Importers and Distributors of High-Risk AI Systems 20
6. Limited-Risk AI 28
7. Generative AI 30
8. Innovation and Regulatory Sandboxes 34
9. Standards, Specifications and Certificates 38
10. Supervision and Enforcement 41

List of Abbreviations

AI Act	*European Union's Artificial Intelligence Act*
Board	*European Union Artificial Intelligence Board*
CAB	*Conformity Assessment Body*
Commission	*European Commission*
GDPR	*European Union General Data Protection Regulation*
GPAI Model	*General-Purpose Artificial Intelligence Model*
NB	*Notified Body*
SME	*Small and Medium-Sized Enterprise*

How We Can Help

WilmerHale has a leading practice in EU law and regulation, advising clients on high-profile matters in both established and emerging market sectors across a wide variety of industries. With around 1,100 lawyers located throughout 12 offices in the United States, Europe and the United Kingdom, we offer a global perspective on EU law issues and provide single-team transatlantic and Europe-wide services. We practice at the very top of the legal profession and offer a cutting-edge blend of capabilities that enables us to handle cases of any size and complexity.

Our European offices in Brussels, Frankfurt, Berlin and London are best known for high-quality regulatory work before national and European authorities and appellate work before EU Courts. Clients entrust us with complex cases because of our expertise, reliability, responsiveness, precision, and reputation with regulators. Our European team is involved in a huge number of cases in various areas of EU law, including several precedent-setting data protection and competition law cases. In addition, many of our lawyers are qualified in several jurisdictions across the EU, its neighbouring countries, and the United States and can handle the most complex cases requiring native-speaker proficiency in multiple languages.

Our European team works seamlessly with our US AI team, leveraging our combined legal expertise to provide comprehensive, cross-border support on AI-related matters. This close collaboration ensures that our clients benefit from globally informed legal strategies.

For more information on this guide or other AI or data-related matters, please contact one of the authors.

Dr. Martin Braun
Partner
Frankfurt/Brussels

Anne Vallery
Partner-in-Charge
Brussels

Itsiq Benizri
Counsel
Brussels

1.

Scope and Approach of the AI Act

Material Scope – What Is AI?

– **AI Systems.** The definition of "AI system" in the AI Act is inspired by the OECD definition, which is widely accepted. It focuses on two key characteristics of AI systems: (1) they operate with varying levels of autonomy and (2) they infer from the input they receive how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments.

Article 3(1) of the AI Act:

"AI system" means a machine-based system designed to operate with varying levels of autonomy, 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.

Recital 12 of the AI Act provides additional background regarding the intentions of the legislators with regard to the definition of AI systems:

[This] definition should be based on key characteristics of AI systems that distinguish it from simpler traditional software systems or programming approaches and should not cover systems that are based on the rules defined solely by natural persons to automatically execute operations. A key characteristic of AI systems is their capability to infer. This capability to infer refers to the process of obtaining the outputs, such as predictions, content, recommendations, or decisions, which can influence physical and virtual environments, and to a capability of AI systems to derive models or algorithms from inputs or data. The techniques that enable inference while building an AI system include machine learning approaches that learn from data how to achieve certain objectives, and logic- and knowledge-based approaches that infer from encoded knowledge or symbolic representation of the task to be solved. The capacity of an AI system to infer transcends basic data processing by enabling learning, reasoning or modelling. The term "machine-based" refers to the fact that AI systems run on machines.

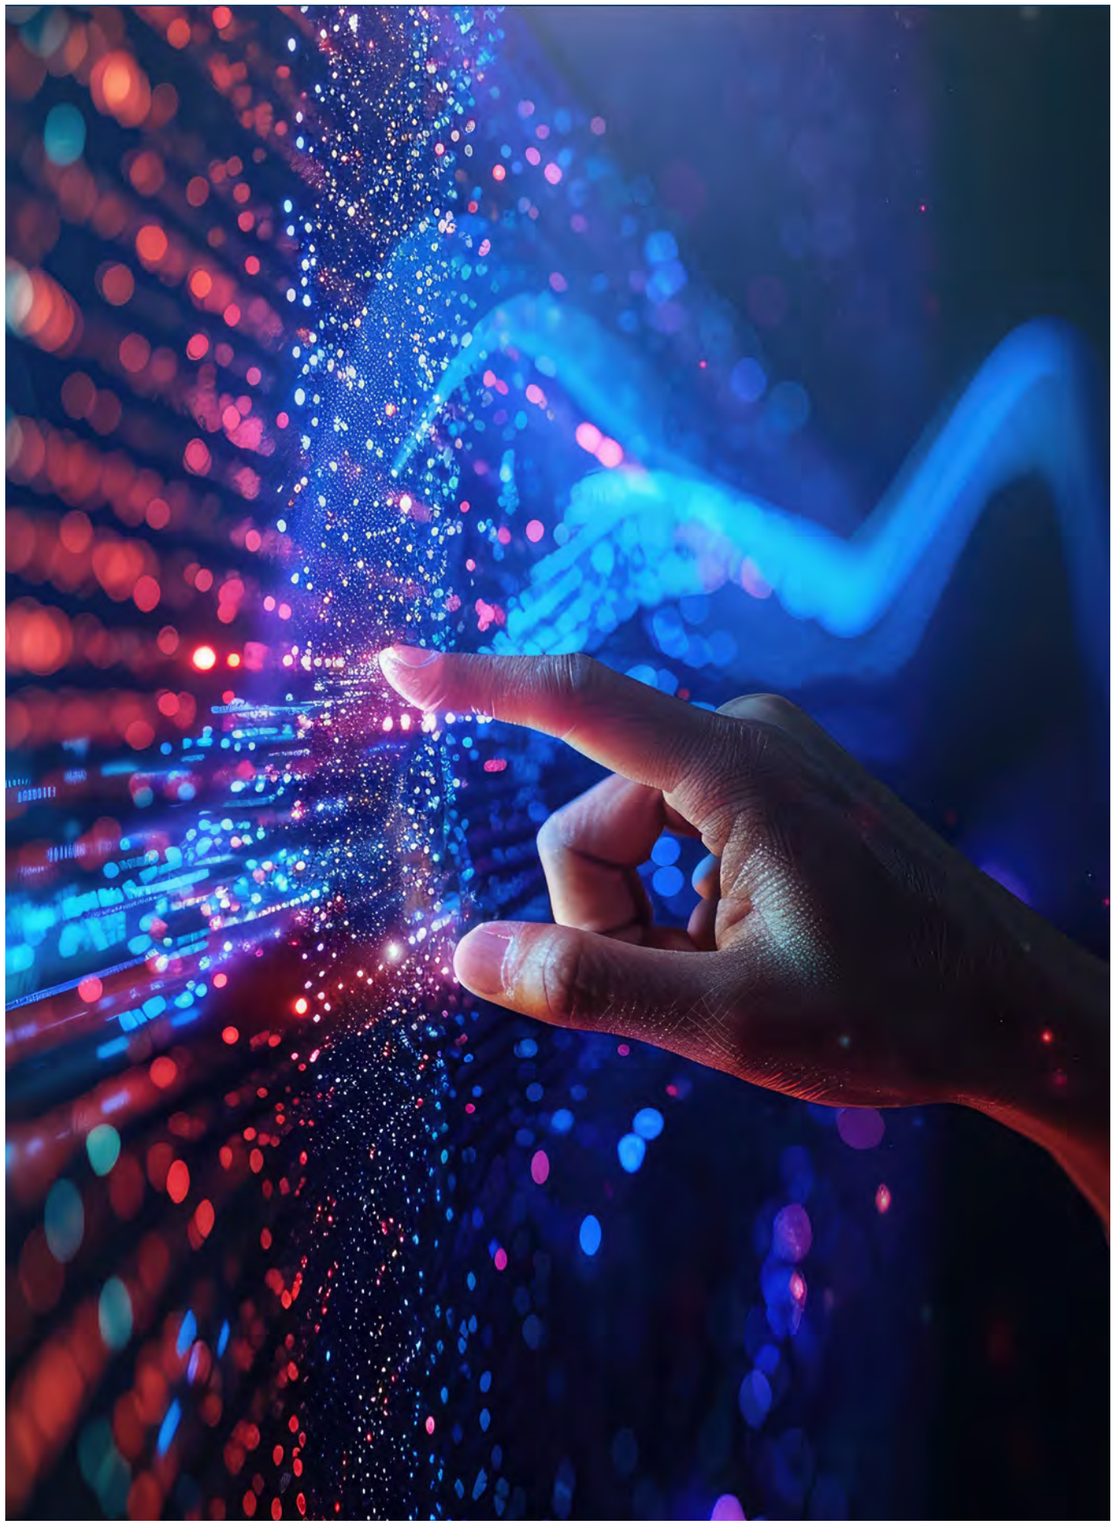

The Commission had also adopted guidelines on the definition of AI systems.

- **General-Purpose AI Models/Generative AI.** During the negotiations, a chapter on general-purpose AI models was added to the AI Act. The legislation now differentiates between "general-purpose AI models" (GPAI Models), a subcategory "general-purpose AI models with systemic risk", and general-purpose AI models with high-impact capabilities.
  - AI models are a component of an AI system and are the engines that drive the functionality of AI systems. AI models require the addition of further components, such as a user interface, to become AI systems.
  - While the AI Act generally does not subject AI models to legal obligations, it defines "GPAI model" as an AI model that (1) displays significant generality; (2) is capable of competently performing a wide range of tasks; and (3) can be integrated into a variety of downstream systems or applications.
  - AI models used for research, development, or prototyping activities before market release are not covered under the AI Act.

Personal Scope – Who Is Subject to the AI Act?

The AI Act identifies and defines the following key players, all of which can be natural or legal persons.

- **Providers** develop or have developed AI systems or GPAI Models with a view to placing them on the market or putting them into service under their own name or trademark, whether for payment or free of charge. The terms "placing on the market" and "putting into service" refer to specific concepts defined in the AI Act:
  - **Placing on the European Union's market.** A company or an individual places an AI system on the market when it *first* makes it available in the European Union.
  - **Putting into service in the European Union.** A provider puts an AI system into service by supplying such a system for *first* use directly to a deployer or for its own use within the European Union for the system's intended purpose.
- **Importers** are located or established in the European Union and place on the market AI systems bearing the name or trademark of a natural or legal person established outside the European Union.
- **Distributors** are players in the supply chain, other than the provider or the importer, that make an AI system available on the EU market.
- **Deployers** use AI under their authority in the course of their professional activities. In practice, it is likely that companies will very quickly be above this very low threshold.

Territorial Scope – Where Does the AI Act Apply?

The AI Act has significant extraterritorial effects, as it applies to providers who place or put into service AI systems on the EU market, irrespective of where they are established or located. The AI Act also applies to providers and deployers established or located outside the EU in cases where the output of the system is used in the EU. The AI Act obviously also applies to deployers who are established or located in the EU. For affected individuals, the AI Act only applies when they are in the EU. There is little clarity or precision regarding distributors.

AI Outside the Scope of the AI Act

The AI Act does not apply to AI specifically developed and put into service for the sole purpose of scientific research and development. The AI Act does not apply to any research, testing or development activity that occurs before an AI system is placed on the market or put into service — but this exemption does not apply to real-world testing. In addition, the AI Act does not apply to systems released under free and open-source licenses, unless such systems qualify as high-risk, prohibited or generative AI. Finally, the AI Act is not applicable to AI systems used ***solely*** for military, defence, or national security purposes, irrespective of the entity performing those activities.

What Is the EU Approach to AI Regulation?

The AI Act relies on a risk-based approach, which means that different requirements apply in accordance with the level of risk.

- **Unacceptable risk (see Chapter 3).** Certain AI practices are considered to be a clear threat to fundamental rights and are prohibited. The respective list in the AI Act includes AI systems that manipulate human behaviour or exploit individuals' vulnerabilities (e.g., age or disability) with the objective or the effect of distorting their behaviour. Other examples of prohibited AI include certain biometric systems, such as emotion recognition systems in the workplace or real-time categorisation of individuals.
- **High risk (see Chapters 4 and 5).** AI systems identified as high-risk will be required to comply with strict requirements, including risk-mitigation systems, high-quality data sets, logging of activity, detailed documentation, clear user information, human oversight, and a high level of robustness, accuracy and cybersecurity. Examples of high-risk AI systems include critical infrastructures, such as energy and transport, medical devices, and systems that determine access to educational institutions or jobs.
- **Limited risk (see Chapter 6).** Providers must ensure that AI systems intended to directly interact with natural persons, such as chatbots, are designed and developed in such a way that individuals are informed that they are interacting with an AI system. Typically, deployers of AI systems that generate or manipulate deepfakes must disclose that the content has been artificially generated or manipulated.
- **Minimal risk.** There are no restrictions on minimal-risk AI systems, such as AI-enabled video games or spam filters. Companies may, however, commit to voluntary codes of conduct.

Relationship With the EU General Data Protection Regulation

EU laws on the protection of personal data, privacy and the confidentiality of communications continue to apply to the processing of personal data in connection with the AI Act. The AI Act does not affect the EU General Data Protection Regulation (GDPR) and the ePrivacy Directive 2002/58/EC.

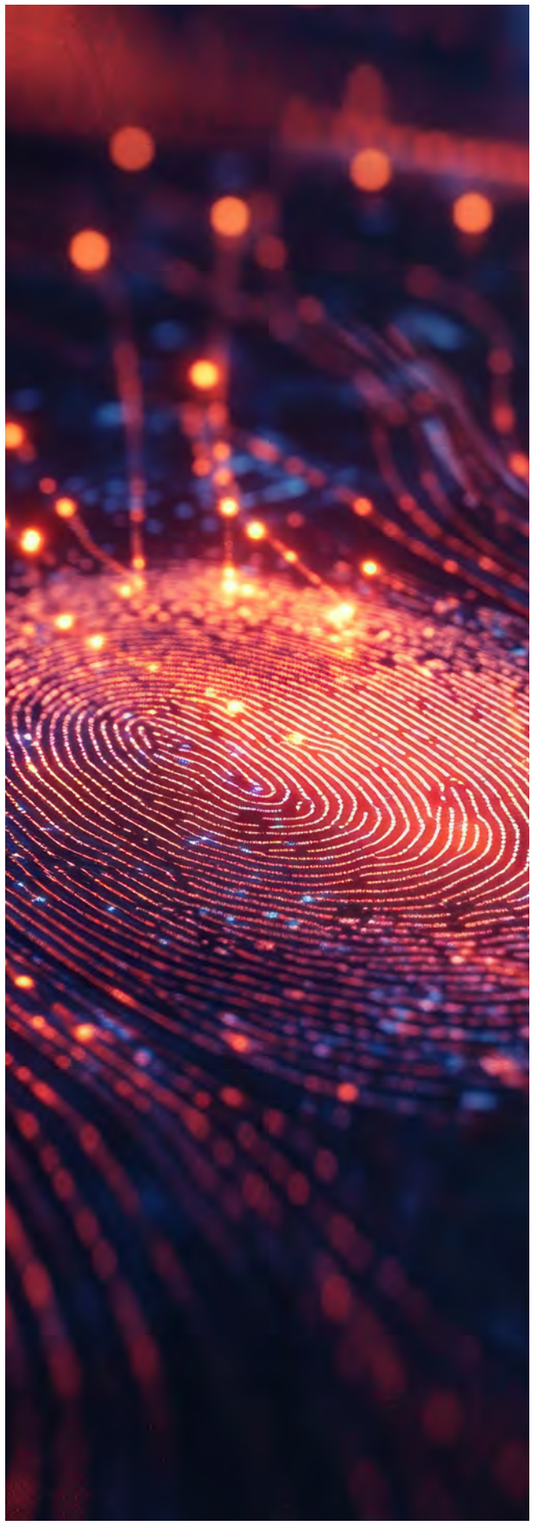

2.

Critical Milestones on the Road to Full Applicability of the AI Act

The AI Act was published in the Official Journal of the European Union on July 12, 2024.

While the AI Act will generally apply starting on August 2, 2026, the exact milestones are quite nuanced and complex, with some provisions already applying since February 2, 2025.

Below, we set out the key dates for the various operators, especially providers and deployers, as well as the dates by which the Commission will have to prepare implementing acts, documentation and reports to help the operators ensure compliance with the AI Act.

2024 August 1

Entry into force of the AI Act (Article 113). This means that the AI Act became part of the EU legal order. It does not mean that the provisions of the AI Act became applicable on that date.

November 2 2024

By this date, Member States had to identify the public authorities or bodies that supervise or enforce obligations under EU law protecting fundamental rights, including the right to nondiscrimination, in relation to the use of high-risk AI systems referred to in Annex III of the AI Act (Article 77(2)). This has not been done in all EU Member States yet.

2025 February 2

Chapters I and II of the AI Act apply from this date (Article 113(a)). These include the Act's general provisions (e.g., geographic scope, definitions) and its provisions on prohibited AI practices.

The general obligation to ensure a sufficient level of AI literacy of staff under Article 4 of the AI Act will also apply from this date.

May 2 2025

By this date, codes of practice for the implementation of general-purpose AI models and related obligations must be ready (Article 56(9)). These codes should support providers in achieving compliance with their duties relating to general-purpose AI models.

2025 August 2

From this date, Chapter III, Section 4 (Notifying authorities and notified bodies), Chapter V (General-purpose AI models), Chapter VII (Governance), and Chapter XII (Penalties) will apply (except for Article 101, which deals with fines for providers of general-purpose AI models).

- Chapter III, Section 4 deals with notifying authorities and notified bodies, which are essential for the establishment of conformity assessment bodies.
- Chapter V contains the provisions related to general-purpose AI models introduced late in the legislative process; for example, the mandatory notification procedure for the provider (Article 52 (1)), documentation requirements (Article 53), and the appointment of an authorised representative (Article 54). Article 55 contains additional responsibilities focusing on the evaluation and mitigation of systemic risk and cyber and infrastructure security.
- Chapter VII sets out the EU's AI-related governance structure, including the AI Office, the European Artificial Intelligence Board, the advisory forum and the scientific panel. On the Member State level, the competent authorities must be appointed by this date (Article 70(2)).

By the same date, the Commission must finalise its guidance to facilitate compliance with the reporting obligations in case of serious incidents (Article 73(7)).

- Chapter XII deals with penalties. This includes Article 99(3), which specifies the fines for noncompliance with prohibited AI practices referred to in Article 5. These fines can reach €35 million, or up to 7% of worldwide annual revenue, if the offender is an undertaking.

February 2 2026

By this date, the Commission must issue implementing acts creating a template for high-risk AI providers' post-market monitoring plans, which should serve as the basis for said monitoring system established by Article 72.

Similarly, the Commission must, by this date, provide guidelines for the practical implementation of Article 6 concerning the classification of an AI system as high risk (Article 6(5)).

2026 August 2

This is the default date by which the provisions of the AI Act become applicable.

The obligations regarding high-risk AI systems will apply from this date, including those related to risk and quality management systems, diligent data governance, technical documentation, recordkeeping, and transparency and clear user information obligations.

Chapter IV addresses operators of AI systems directly interacting with humans, generative AI systems, and emotion recognition or biometric categorisation systems, introducing disclosure and information responsibilities.

By this date, Member States must have implemented rules on penalties and other enforcement measures and notified the Commission about them (Article 99).

Member States must have established at least one AI regulatory sandbox, which must be operational at a national level (Article 57(1)).

August 2 2027

This is the ultimate deadline for AI systems covered by existing harmonisation legislation (Article 113(c)) and for providers of general-purpose AI models that have been placed on the market for up to 12 months after August 1, 2024, to comply with the AI Act.

3.

Prohibited AI Practices

Article 5 of the AI Act essentially prohibits AI practices that materially distort people's behaviour or that involve discrimination, profiling or other practices that raise serious concerns in democratic societies. Some prohibitions are particularly relevant from a business perspective. Others are most likely to be relevant for governments or only apply in the context of law enforcement.

The list of prohibited AI practices is not set in stone. The Commission will assess the need to amend this list once a year and share its findings with European Union lawmakers (Article 112). The Commission has already adopted guidelines for the practical implementation of the AI Act provisions regarding prohibited AI systems.

AI Systems Prohibited for Businesses

The AI Act prohibits placing AI systems on the European Union's market, putting them into service, or using them in the European Union to materially distort people's behaviour in a manner that causes or is likely to cause them physical or psychological harm:

- **Prohibited Practices.** The AI Act prohibits placing on the market, putting into service, and using certain AI systems. There is no specific definition for "use of AI" in the AI Act, which suggests a common and broad understanding of the term.

- **Prohibited Systems.** The AI Act prohibits placing on the market, putting into service, and using the following AI systems:
  - **Subliminal, manipulative and deceptive systems.** These are AI systems that deploy subliminal techniques beyond a person's consciousness or purposefully use manipulative or deceptive techniques that materially distort people's behaviour by appreciably impairing their ability to make informed decisions. Such systems cause people to make decisions that they would not have otherwise taken, [likely] resulting in significant harm.
  - **Systems that exploit vulnerabilities.** These are AI systems that exploit people's vulnerabilities due to their age, disability, or social or economic situation. Such systems also distort people's behaviour, [likely] resulting in significant harm.

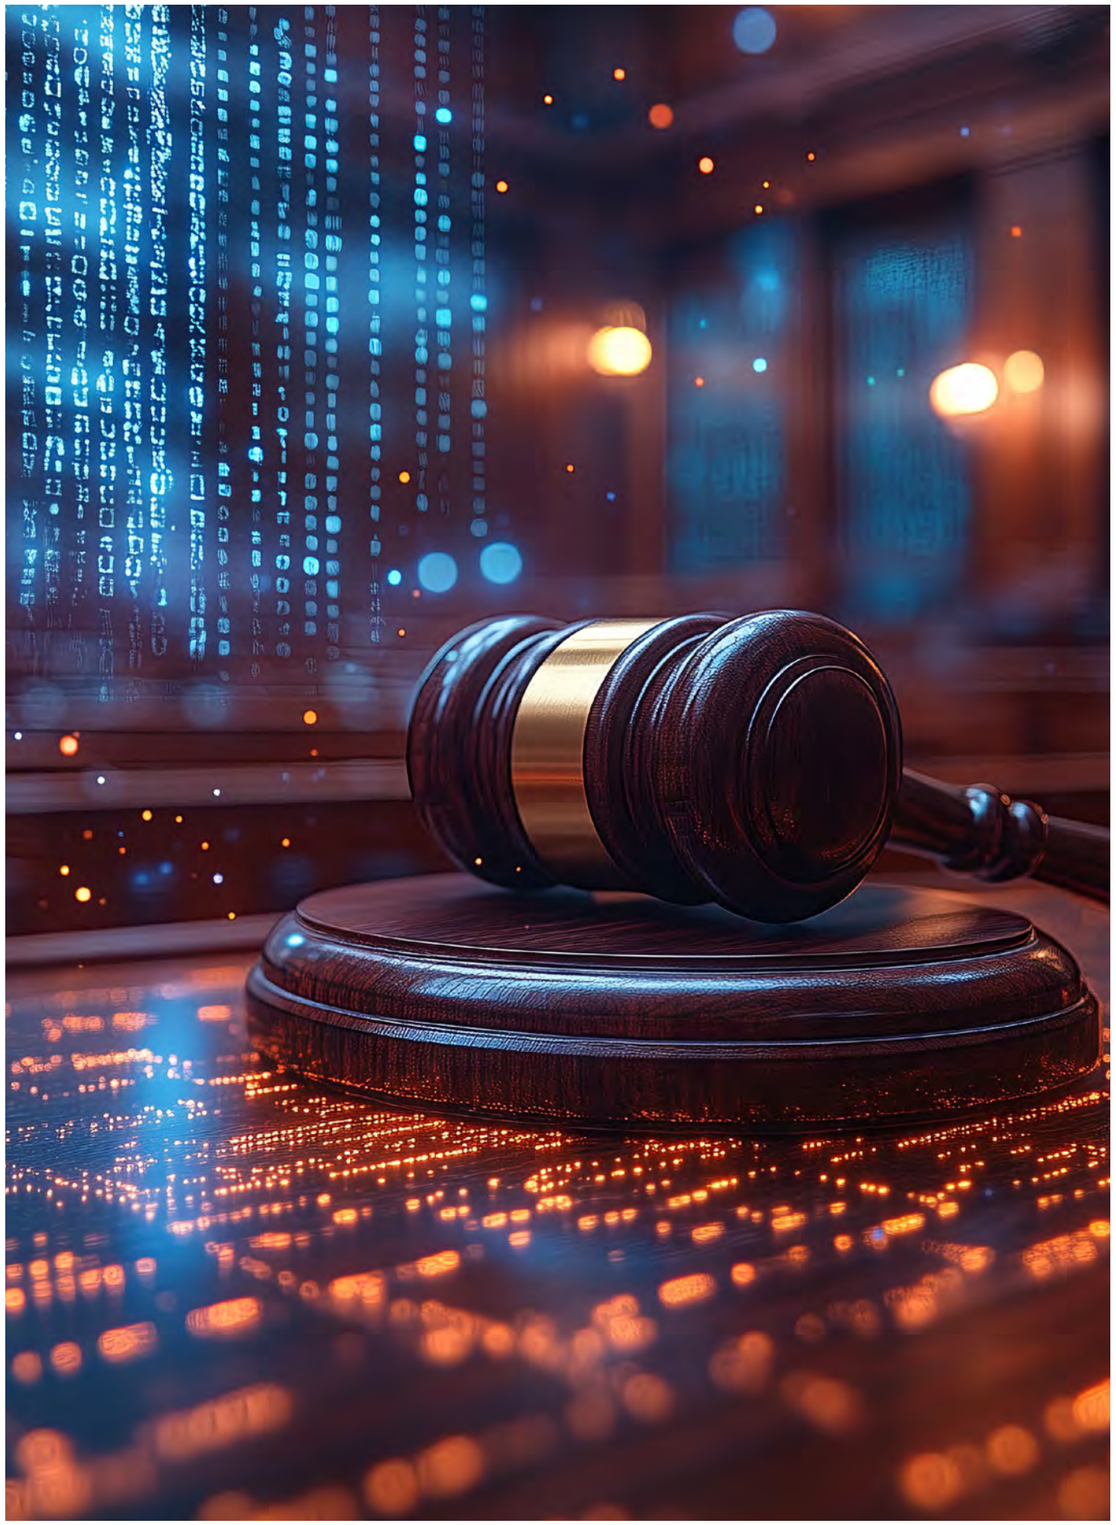

- **Facial recognition databases.** These are AI systems that create or expand facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage.
- **Systems that infer emotions.** These are AI systems that infer emotions of individuals in the workplace and educational institutions, except for AI medical or safety systems.
- **Systems using biometric categorisation.** These are AI systems that categorise individual natural persons based on their biometric data to deduce or infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation. Importantly, the processing of biometric data for the purpose of uniquely identifying an individual is subject to strict restrictions under the GDPR. Such processing is prohibited unless one of the limited exceptions applies, such as the data subject's explicit consent.

Prohibited Systems Likely to Be Used by Governments

The AI Act prohibits placing AI systems on the European Union's market, putting them into service or using them in the European Union for social scoring or "minority report" purposes.

- **Social scoring.** This refers to AI systems used for the evaluation or classification of people based on their social behaviour or known, inferred, or predicted personal characteristics. The prohibition applies where such social scoring leads to a detriment or unfavourable treatment:
  - in social contexts that are unrelated to the contexts in which the data was originally generated or collected; and/or
  - the detriment or unfavourable treatment is unjustified or disproportionate to the social behaviour in question or its gravity.
- **Minority report.** This refers to AI systems used to make risk assessments of individuals to identify or predict the risk that they will commit a criminal offense based solely on their profiling or on an assessment of their personality traits and characteristics. This prohibition, however, does not apply to AI systems used to support the human assessment of the involvement of a person in a criminal activity, which is already based on objective and verifiable facts.

AI Systems Prohibited for Law Enforcement Purposes

The AI Act prohibits the use of real-time remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement. The AI Act does not prohibit placing on the market or putting such systems into service. The prohibition applies unless and in addition to specific safeguards, in as far as the use of real-time remote biometric identification is strictly necessary for:

- **the targeted search for specific victims** of abduction, human trafficking, or sexual exploitation, and the search for missing persons;
- **the prevention of a specific, substantial and imminent threat** to the life or physical safety of natural persons, or a genuine and present or foreseeable threat of a terrorist attack; or
- **the localisation or identification of persons suspected of having committed a criminal offense,** for the purposes of conducting a criminal investigation or prosecution or executing a criminal penalty. This only applies to specific offenses listed in the AI Act and punishable by a custodial sentence or a detention order for a maximum period of at least four years.

4. Definition and Requirements for High-Risk AI Systems

In this section, we will focus on the identification of "high-risk AI systems" under the AI Act and the requirements that apply to such systems.

Identifying High-Risk AI Systems

Article 6 of the AI Act describes the thresholds that lead to an AI system being "high risk." Either such system meets the criteria in Article 6(1) of the AI Act or it falls into a category referred to in Article 6(2) of the AI Act.

Article 6(1) of the AI Act. An AI system will be considered high risk if two cumulative conditions are fulfilled:

1. The AI system is intended to be used as a safety component of a product (or is a product) covered by specific EU harmonisation legislation listed in Annex I of the AI Act. This list contains more than 30 directives and regulations, including legislation regarding the safety of toys, vehicles, civil aviation, lifts, radio equipment and medical devices; and
2. The same harmonisation legislation mandates that the product that incorporates the AI system as a safety component, or the AI system itself as a stand-alone product, undergo a third-party conformity assessment before being placed on the EU market or put into service within the EU.

Article 6(2) of the AI Act—Specific List. In addition, the AI Act contains, in its Annex III, a list of AI systems that must be considered high risk. This list currently contains AI systems in eight different categories. Examples include, subject to specific conditions and exemptions, biometrics, critical infrastructures, education and vocational training, employment, worker management, and access to self-employment. The Commission has the power to amend this list.

The AI systems identified in Annex III will not be considered high risk if they do not pose a significant risk of harm to individuals' health, safety or fundamental rights, including by not materially influencing the outcome of decision-making. This exemption applies where one of the following conditions is met:

- the AI system is intended to perform a narrow procedural task;
- the AI system is intended to improve the result of a previously completed human activity;
- the AI system is intended to detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment, without proper human review; or
- the AI system is intended to perform a preparatory task to an assessment relevant for the purposes of the use cases that are listed as high risk.

However, the exemption never applies if the AI system performs profiling of natural persons. Profiling is defined by reference to Article 4(4) GDPR as any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements.

If a provider considers that an AI system benefits from the exemption, it must document its assessment before placing that system or putting it into service in the European Union. The provider must also register the system in an EU database for high-risk AI systems set up and maintained by the Commission.

The Commission will provide guidelines by February 2, 2026, to specify the practical implementation of classification rules for high-risk AI systems, including the conditions for exceptions.

Requirements for High-Risk AI Systems

High-risk AI systems must comply with a significant number of requirements that consider their intended purposes, the generally acknowledged state of the art, and the risk management system put in place. The applicable requirements are as follows:

- **Risk management.** High-risk AI systems require a risk management system running throughout the entire life cycle of the system. The objective is to identify foreseeable risks to health, safety or fundamental rights when the system is used in accordance with its intended purpose and to adopt appropriate and targeted measures to address those risks; to estimate and evaluate the risks that may emerge when the system is used in accordance with its intended purpose, and under conditions of reasonably foreseeable misuse; and to evaluate other risks that may arise based on a post-market monitoring analysis. Importantly, the risk management system concerns only risks that may be reasonably mitigated or eliminated through the development or design of the high-risk AI system or the provision of adequate technical information.
- **Data and data governance.** The training, validation and testing data used to develop high-risk AI systems must be subject to appropriate data governance and management practices appropriate for the intended purpose of the system.
  - **Examples include** relevant design choices; appropriate data collection processes; relevant data preparation processing operations, such as annotation, labelling, cleaning, updating, enrichment and aggregation; the formulation of relevant assumptions; prior assessment of the availability, quantity and suitability of the datasets needed; examination in view of possible biases likely to affect individuals' health and safety, negatively impact fundamental rights, or lead to discrimination prohibited under EU law; appropriate measures to detect, prevent and mitigate those biases; and identification of relevant data gaps or shortcomings that prevent compliance, and how they can be addressed.
  - **Training, validation and testing datasets** must be relevant, sufficiently representative, and to the best extent possible free of errors and complete in view of their intended purpose. They must have the appropriate statistical properties, including, where applicable, as regards the persons or groups of persons in relation to whom the high-risk AI system is intended to be used. Those characteristics of the datasets may be met at the level of individual datasets or at the level of combinations of datasets. In addition, the datasets must consider, to the extent required by the intended purpose, the characteristics or elements that are particular to the specific geographical, contextual, behavioural or functional setting within which the AI system is intended to be used.

- **For AI systems that are not developed based on AI model training**, those requirements apply only to the testing datasets.

— **Technical documentation.** Technical documentation for high-risk AI systems must be drawn up before the system is placed or put into service in the European Union. Such documentation must demonstrate that the system complies with the requirements set out in the AI Act.

- **The AI Act provides a list of the minimum information that the technical documentation must include**, such as a description of the system, its elements and the process for its development; information about the monitoring, functioning and control of the system; a description of the appropriateness of the performance metrics for the system; a description of the risk management system; a record of the relevant changes made by the provider through the life cycle of the system; the technical standards applied; the declaration of conformity; and the system in place to evaluate the system performance.
- **SMEs, including startups**, may provide the elements of the technical documentation in a simplified manner. The Commission will publish a simplified form to that end.

— **Recordkeeping.** High-risk AI systems must allow for the automatic recording of events (logs) over their lifetime. The objective is to ensure the traceability of the functioning of the system to ensure that it is appropriate to its intended purpose. To that end, logging capabilities must enable the recording of events relevant for identifying situations that may result in the system presenting a substantial modification or that have the potential to adversely affect individuals' health, safety or fundamental rights to a degree that goes beyond that considered reasonable and acceptable in relation to its intended purpose, or under normal or reasonably foreseeable conditions of use. These logging capabilities must also facilitate post-market monitoring; and monitoring of the operation of the systems deployed by financial institutions.

— **Transparency and provision of information to deployers.** Deployers must be provided with sufficiently transparent information to interpret the system's output and use it appropriately. The system must be accompanied by instructions for use in an appropriate format that includes concise, correct and clear information that is relevant, accessible and comprehensible. The instructions for use must contain at least the following information: the providers' identity and contact details; the system characteristics, capabilities and limitations of performance; changes to the system and its performance; human oversight measures; the computational and hardware resources needed; and, where relevant, the mechanisms included within the system that allows users to properly collect, store and interpret the logs.

— **Human oversight.** High-risk AI systems must be designed and developed in such a way that they can be effectively overseen by humans. Human oversight must aim to prevent or minimise the risks to health, safety or fundamental rights that may emerge when a high-risk AI system is used in accordance with its intended purpose or under conditions of reasonably foreseeable misuse. The oversight measures must be commensurate to the risks, level of autonomy and context of use.

- **Human oversight must be achieved through at least one of the following types of measures:**
  - measures identified and built, when technically feasible, into the system by the provider before it is placed on the EU's market or put into service in the EU; or
  - measures that are identified by the provider before placing the system on the market or putting it into service in the EU and that are appropriate to be implemented by the deployer.
- **Individuals to whom oversight is assigned must be able, as appropriate and proportionate to the circumstances, to:**
  - properly understand the relevant capacities and limitations of the system and monitor its operations, including detecting and addressing anomalies, dysfunctions and unexpected performance;
  - remain aware of the possible tendency of automatically relying or over-relying on the output produced by the system;
  - correctly interpret the system's output;
  - decide not to use the system or otherwise disregard, override or reverse the system's output; and
  - intervene in the operation of the system or interrupt it through a "stop" button or a similar procedure that allows the system to come to a halt in a safe state.

— **Accuracy, robustness and cybersecurity.** High-risk AI systems must be designed and developed in such a way that they achieve an appropriate level of accuracy, robustness and cybersecurity and perform consistently in those respects throughout their life cycle. The Commission will encourage the development of benchmarks and measurement methodologies to that effect.

- **The levels of accuracy and the relevant accuracy metrics** must be declared in the instructions of use.
- **High-risk systems must be as resilient as possible** regarding errors, faults or inconsistencies that may occur within the system or the environment in which it operates.
- **High-risk AI systems that continue to learn after being placed on the market or put into service** must be developed in such a way as to eliminate or reduce as far as possible the risk of possibly biased outputs influencing input for future operations, and to ensure that any such feedback loops are duly addressed with appropriate mitigation measures.
- **High-risk AI systems must be resilient against attempts by unauthorised third parties** to alter their use, outputs or performance by exploiting system vulnerabilities, and the technical solutions aiming to ensure the cybersecurity of high-risk AI systems must be appropriate to the risks and circumstances.

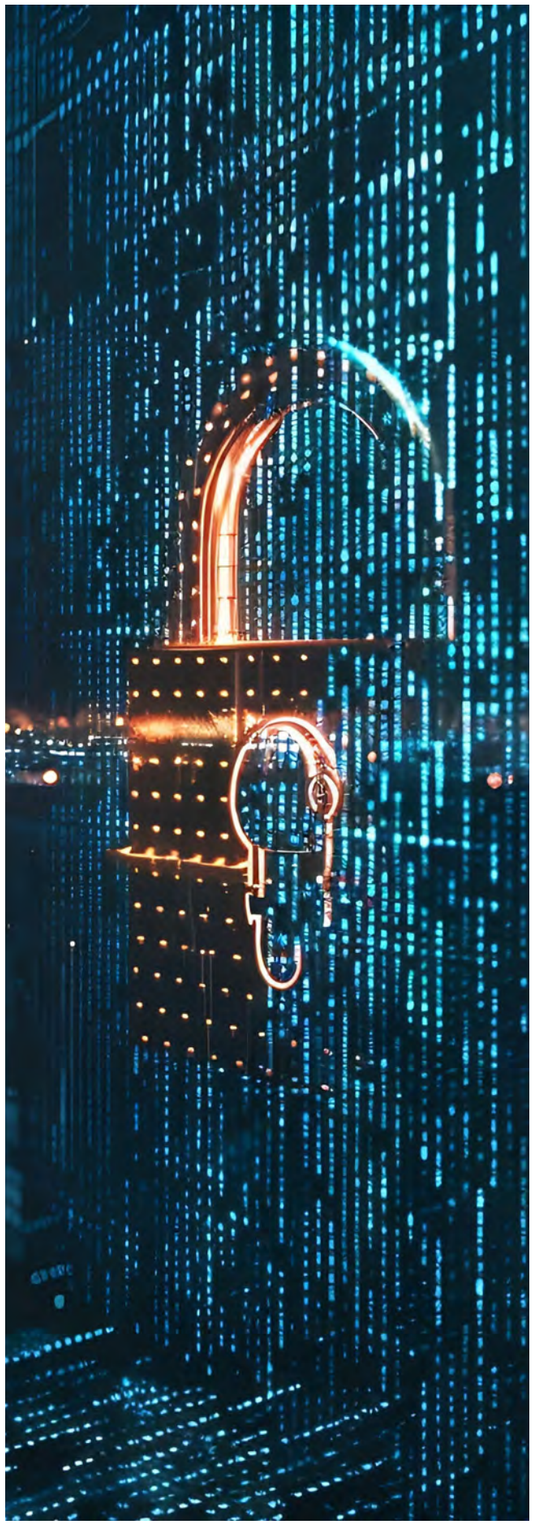

5.

Obligations for Deployers, Providers, Importers and Distributors of High-Risk AI Systems

This chapter focuses on obligations that the AI Act sets for deployers, providers, importers and distributors regarding high-risk AI systems.

Obligations for Deployers of High-Risk AI Systems

- **Instructions for use.** Deployers must take appropriate technical and organisational measures to ensure they use high-risk AI systems in accordance with the instructions for use. EU or national law can impose additional obligations in this respect.
  - **Monitoring.** Deployers must monitor the operation of the system based on the instructions for use. Where relevant, deployers must inform providers.
  - **Risk to health, safety or fundamental rights.** Where deployers have reasons to believe that using the system in accordance with the instructions may adversely affect individuals' health, safety or fundamental rights (see above), they must, without undue delay, inform the provider or distributor and the relevant market surveillance authority. They should also suspend the use of the system.
  - **Serious incident.** Where deployers have identified a serious incident, they must immediately inform first the provider and then the importer or distributor and the relevant market surveillance authorities. If the deployer is unable to contact the provider, it must inform the market surveillance authority of the European country where the incident occurred. This should occur immediately after the deployer establishes a causal link between the AI system and the serious incident, or the reasonable likelihood of such a link. In any case, this notification should take place no later than 15 days after the deployer becomes aware of the incident.
  - **Logs.** Deployers of high-risk AI systems must retain the logs automatically generated by the system, to the extent that such logs are within their control, for a duration appropriate to the system's intended purpose but of at least six months, unless provided otherwise in applicable EU or national law.

- **Input data.** If the deployer exercises control over the input data, it must ensure that such data is relevant and sufficiently representative in view of the intended purpose of the system.

— **Human oversight.** Deployers must assign human oversight to individuals who have the necessary competence, training, authority and support. Deployers are free to organise their own resources and activities to implement the oversight measures indicated by the provider. EU or national law can impose additional obligations. The above requirement regarding input data also applies.

— **Workplace.** Before putting into service or using a high-risk AI system in the workplace, deployers that are employers must inform workers' representatives and the affected workers that they will be subject to the use of a high-risk AI system.

— **Transparency.** Deployers of specific high-risk AI systems listed in the AI Act (e.g., those used in critical infrastructures, education and vocational training, employment, worker management, and access to self-employment) that make decisions or assist in making decisions related to natural persons must inform these persons that they are subject to the use of the high-risk AI system.

— **Cooperation with authorities.** Deployers must cooperate with the relevant national competent authorities in any action those authorities take in relation to the high-risk AI system to implement the AI Act.

— **Fundamental rights impact assessment.** Before deploying high-risk AI systems to evaluate individuals' creditworthiness, establish their credit score (excluding systems used to detect financial fraud), or assess risks and determine pricing for life and health insurance, deployers must assess the impact on fundamental rights that the use of such system may entail. This assessment must consider the processes in which the system will be employed, the duration and frequency of its usage, the categories of individuals affected, the specific risks of harm, the measures for human oversight, and the actions to be taken if risks materialise.

- **First use.** This obligation only applies to the first use of the high-risk AI system. The deployer may, in similar cases, rely on previous fundamental rights impact assessments or existing assessments carried out by the provider. However, the deployer needs to update such assessments as appropriate.
- **Notification to authorities.** The deployer must inform the market surveillance authority of the results of its assessment, with only very limited exemptions.
- **Data protection impact assessment.** If any of the obligations in relation to the fundamental rights impact assessment are already complied with as a result of a GDPR data protection impact assessment, the fundamental rights impact assessment must complement that data protection impact assessment.

Obligations for Providers of High-Risk AI Systems

Providers of high-risk AI systems must ensure that their systems comply with the requirements associated with such systems and demonstrate such compliance to national competent authorities upon request. Providers must also indicate on their system or, if that is not possible, on the packaging or accompanying documentation, their name, registered trade name or trademark, and the address at which they can be contacted. In addition, providers must comply with the following requirements.

— **Put in place a quality management system to ensure compliance.** This system must be documented in a systematic and orderly manner, comprising written policies, procedures and instructions, in proportion to the size of the provider. The system must include minimum information as listed in the AI Act, such as a strategy for regulatory compliance; techniques, procedures, and systematic actions for the design control and verifications; examination, test and validation procedures during and after the development of the system; technical standards and specifications to ensure compliance; risk management and post-market monitoring systems; incident reporting procedures; and an accountability framework setting out individuals' responsibilities.

— **Keep the required documentation for 10 years after the system has been placed on the market or put into service in the European Union.** This documentation must include the technical documentation and the documentation concerning the quality management system (see **chapter 4** for more details), the EU declaration of conformity, and any document issued by conformity assessment bodies.

— **Keep the logs automatically generated by the system to the extent they are under providers' control.** Providers must keep the logs for a period appropriate to the intended purpose of the system but of at least six months, unless provided otherwise in relevant EU or national law.

— **Ensure that the system undergoes the conformity assessment procedure before being placed on the market or put into service in the European Union.** This procedure varies depending on the type of high-risk system. Providers of AI systems used for biometric purposes can choose either an internal control procedure or an external control by a conformity assessment body, provided they have applied specific technical standards. For other high-risk AI systems identified in the AI Act, providers can follow the conformity assessment procedure based on internal control. Specific rules apply to AI systems covered by EU harmonised legislation. Essentially, for some of them, the main rule is that providers must follow the procedure required under the relevant legislation.

— **Draw up an EU declaration of conformity with the requirements associated with high-risk AI systems.** The provider must draw up a written, machine-readable physical or electronically signed EU declaration of conformity for each high-risk AI system and keep it at the disposal of the national competent authorities for 10 years after the system has been placed on the market or put into service. The declaration of conformity must contain the information set out in the AI Act. The Commission may update this list in future. This information includes, for example, information allowing the identification and traceability of the system, a statement that the declaration of conformity is issued under the sole responsibility of the provider, and references to technical standards or specifications in relation to which conformity is declared.

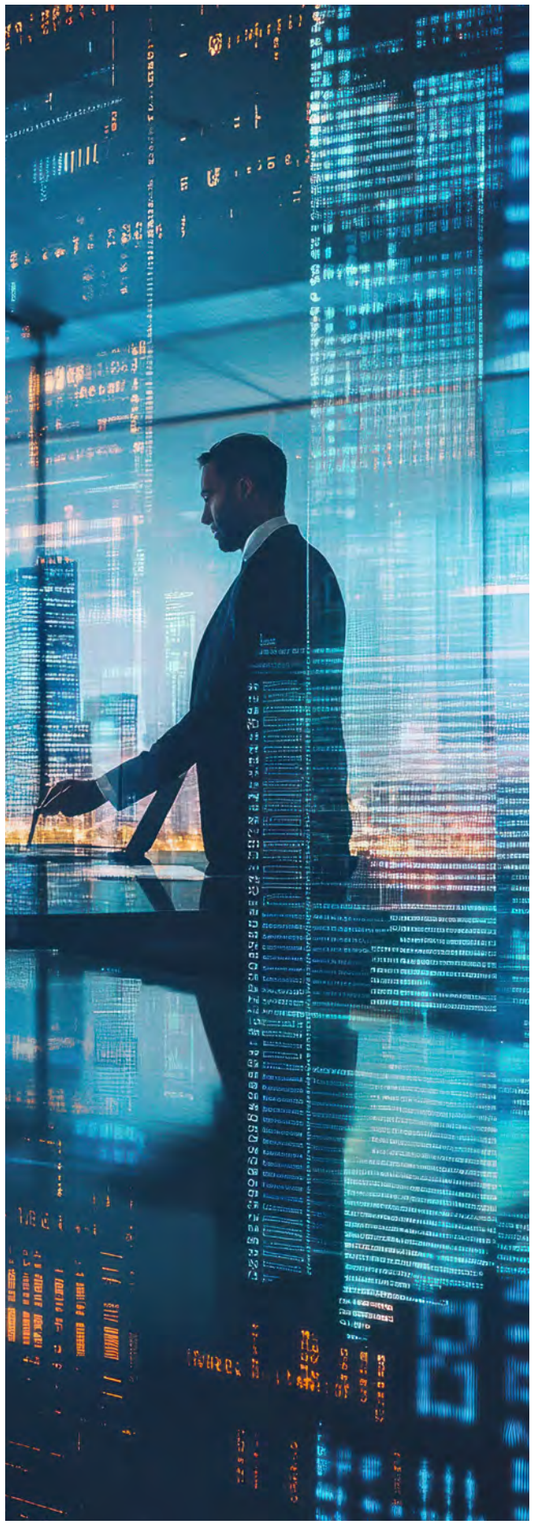

— **Affix the CE marking to the system** or, where that is not possible, on its packaging or its accompanying documentation to indicate conformity with the AI Act. The marking refers to the letters "CE", signifying that a product sold in the European Union has been assessed to meet the relevant protection requirements.

— **Comply with the registration obligations.** Before placing a high-risk AI system on the market or putting it into service (except for critical infrastructures), providers (or authorised representatives) must register themselves and their systems in the EU database (see **chapter 4** for more details).

— **Ensure post-market monitoring.** Providers must establish and document a post-market monitoring system in a manner that is proportionate to the nature of the AI technologies and the risks of the high-risk AI system. The monitoring system must actively and systematically collect, document and analyse relevant data on the performance of high-risk AI systems throughout their lifetime that may be provided by deployers or collected through other sources, and that allow the provider to evaluate the continuous compliance of AI systems with the AI Act. The post-market monitoring system must be based on a post-market monitoring plan, which should be part of the technical documentation drawn up before the AI system is placed on the market or put into service in the European Union. The Commission will create a template for the monitoring plan and specify the elements that it should include.

— **Report serious incidents.** Providers must report any serious incident to the market surveillance authorities of the European country where that incident occurred. Serious incidents are incidents or the malfunctioning of an AI system that (in)directly leads to the death of a person or serious harm to a person's health, serious and irreversible disruption of the management or operation of critical infrastructure, infringement of obligations under EU law intended to protect fundamental rights, or serious harm to property or the environment. In specific cases, the reporting requirement is limited to the two latter cases.

The timing for reporting serious incidents varies depending on the context. Where necessary to ensure timely reporting, providers or, where applicable, deployers may submit an initial incomplete report followed by a complete one.

- **General rule.** In general, providers must report serious incidents immediately after having established a causal link between the AI system and the incident, or the reasonable likelihood of such a link. In any event, taking into account the severity of the incident, providers must make the report no later than 15 days after they or, where applicable, deployers become aware of the incident.
- **Critical infrastructures and widespread infringement.** The report must be provided immediately and not later than two days after the provider or, where applicable, the deployer becomes aware of an incident or malfunctioning of an AI system that leads to a serious and irreversible disruption of the management or operation of critical infrastructure, or of a widespread infringement. A widespread infringement consists of any act or omission that is contrary to EU law protecting the interests of individuals and has harmed or is likely to harm the collective interests of individuals residing in at least two European countries other than the one in which the act or omission originated or took place, the provider (or its authorised representative) is located or established, or the deployer that committed the infringement is established. A widespread infringement may also consist of any acts or omissions contrary to EU law protecting the interests of individuals that have caused, are causing or are likely to cause harm to the collective interests of individuals and have common features, including the same unlawful practice or the same interest being infringed, and are occurring concurrently, committed by the same player, in at least three European countries.
- **Death.** In the event of the death of a person, the report shall be provided immediately after the provider or the deployer has established, or as soon as it suspects, a causal relationship between the high-risk AI system and the serious incident but not later than 10 days after the date on which the provider or, where applicable, the deployer becomes aware of the serious incident.

— **Ensure follow-up on reporting serious incidents.** Following the reporting of a serious incident, the provider must, without delay, perform the necessary investigation, conduct a risk assessment and take corrective action. The provider must also cooperate with the competent authorities (and the conformity assessment bodies, if applicable). In this context, the provider must inform authorities before altering the AI system in a way that may affect any subsequent evaluation of the causes of the incident.

— **Take the necessary corrective actions and provide the required information.** If providers consider that a high-risk AI system is not in conformity with the AI Act, they must immediately take corrective actions to bring that system into conformity, withdraw it, disable it or recall it, as appropriate. Providers must inform distributors and, where applicable, the authorised representative and importers accordingly. Providers must also immediately investigate the causes and inform market surveillance authorities (and possibly conformity assessment bodies) if they become aware of the fact that a high-risk AI system has the potential to adversely affect individuals' health, safety or fundamental rights to a degree that goes beyond that considered reasonable and acceptable in relation to its intended purpose or under normal or reasonably foreseeable conditions of use. In particular, providers must highlight the nature of the noncompliance and of any relevant corrective action taken.

— **Ensure that the high-risk AI system complies with accessibility requirements for certain products and services.** For businesses, this essentially refers to products and services identified in Directive 2019/882. Examples include computers and operating systems for those computers, payment terminals, terminals used for electronic communication or audiovisual media services, and e-readers.

— **Cooperate with competent authorities.** Upon a national authority's reasoned request, providers must supply all the information and documentation necessary to demonstrate the conformity of the high-risk AI system with the AI Act. Upon reasoned request, providers must also give the authority access to the logs automatically generated by the system to the extent they are under the provider's control.

- **Appoint authorised representatives.** Prior to making high-risk AI systems available on the EU market, providers established outside the European Union must appoint an authorised representative established in the European Union. This representative can be addressed, in addition to or instead of the provider, by the competent authorities on all compliance issues. The authorised representative must perform the tasks specified in the written mandate received from the provider. This mandate must empower the representative to carry out the following tasks:
  - Verify that the provider has drawn up the EU declaration of conformity and the technical documentation and has carried out an appropriate conformity assessment procedure.
  - Keep at disposal of the national competent authorities, for 10 years after the system has been placed on the market or put into service, the contact details of the provider, a copy of the EU declaration of conformity, the technical documentation and, if applicable, the certificate issued by the conformity assessment body.
  - Provide the national competent authority, upon reasoned request, with the requested information and documentation necessary to demonstrate conformity with the requirements for high-risk AI systems set out in the AI Act, including access to the logs automatically generated by the system, provided they are under the provider's control.
  - Cooperate with national competent authorities, upon reasoned request.
  - Comply with the registration obligations (see **chapter 4** for more details)—if the registration is carried out by the provider, the authorised representative must ensure that the registration includes the right information.
- **Understand responsibilities along the value chain.** The provider of a high-risk AI system and the third party that supplies such a system or the tools, services, components or processes used or integrated in such a system must, through a written agreement, specify the necessary information, capabilities, technical access and other assistance based on the generally acknowledged state of the art. The objective is to enable the provider to comply with its obligations. However, this requirement does not extend to third parties offering tools, services, processes or components to the public, excluding general-purpose AI models, under a free and open license.
- **Beware of the requalification clause—deployers and others may become providers.** The AI Act incorporates a requalification clause for high-risk AI systems, wherein any third party, such as a distributor, importer or deployer, is requalified as a provider and consequently subjected to the obligations of the provider if they engage in certain actions.
  - **In general.** These actions are as follows: putting their name or trademark on a system already placed on the market or put into service in the European Union; making substantial modifications to such a system that maintains its high-risk status; or modifying its intended purpose in a manner that renders it high-risk. In such cases, the initial provider is no longer the provider. Instead, it must cooperate with the new one, make available the necessary information, and provide the reasonably expected technical access and other assistance required for the fulfilment of the obligations set out in the AI Act. This is without prejudice to the need to observe and protect intellectual property rights, confidential business information, and trade secrets in accordance with EU and national law. If the initial provider clearly specified that its AI system is not to be changed into a high-risk system, there is no obligation to hand over the documentation.
  - **Specific harmonisation legislation.** For high-risk AI systems that are safety components of products covered by specific EU harmonisation legislation listed in the AI Act (e.g., regarding the safety of toys, lifts, radio equipment or medical devices), two actions requalify third parties as providers: the system is placed on the market together with the product under the manufacturer's name or trademark; or the system is put into service under the product manufacturer's name or trademark after the product has been placed on the market.

Obligations for Importers of High-Risk AI Systems

- **Perform certifications.** Importers are required to make several verifications before placing a high-risk AI system on the market. They must ensure that the provider has carried out a conformity assessment, drawn up the required technical documentation, affixed the required CE marking, provided the EU declaration of conformity and instructions for use, and appointed an authorised representative if applicable.
- **Conclude from checks.** If verifications give the importer sufficient reasons to consider that the system is not AI Act–compliant, is falsified, or is accompanied by falsified documentation, the importer cannot place the system on the EU market until it is brought into conformity. Importers must inform the provider, the authorised representative, and the market surveillance authorities if the system in question has the potential to adversely affect individuals' health, safety or fundamental rights to a degree that goes beyond what is considered reasonable and acceptable in relation to its intended purpose or under normal or reasonably foreseeable conditions of use.
- **Be transparent.** Importers must indicate their name, registered trade name or trademark, and address on the system packaging or accompanying documentation, where applicable.
- **Ensure compliance.** Importers are responsible for ensuring that storage or transport conditions do not compromise the system's compliance with the requirements for high-risk AI systems, as detailed in **chapter 4**. This obligation only applies where applicable and while the system is under the importer's responsibility.
- **Keep documentation.** For a period of 10 years after the system has been placed on the market or put into service, importers must keep a copy of the certificate issued by the conformity assessment body, and, where applicable, of the EU declaration of conformity and instructions for use.
- **Cooperate with authorities.** Upon a reasoned request, importers must provide to national competent authorities all the necessary information and documentation to demonstrate the conformity of the system with the AI Act requirements. Importers must also cooperate in any action those authorities take, in particular to reduce and mitigate the risks posed by the system.

Obligations for Distributors of High-Risk AI Systems

- **Perform verifications.** Distributors are required to make different verifications before placing a high-risk AI system on the market. They must ensure that the provider has affixed the required CE marking and provided the EU declaration of conformity and instructions for use. In addition, distributors must ensure that the provider and the importer (as applicable) have complied with their obligation to indicate on the system packaging or accompanying documentation their name, registered trade name or trademark, and address. Distributors must also ensure that providers have put in place an appropriate quality management system (see **chapter 4** for more details).
- **Conclude from checks.** Based on the information available, if a distributor has grounds to believe that the system does not comply with the requirements of the AI Act, it is subject to the same obligations as importers, as outlined above. If the distributor has already made the system available on the market, it must take corrective actions necessary to bring the system into conformity with the AI Act's requirements, including withdrawal or recall. Alternatively, the distributor must ensure that the provider, importer or any relevant operator takes these corrective actions. In cases where the high-risk AI system may adversely affect individuals' health, safety, or fundamental rights (see above), the distributor must immediately inform the provider or importer and the relevant national competent authorities. This notification should include details of the noncompliance and any corrective actions taken.
- **Ensure compliance.** The same obligations apply to distributors as apply to importers regarding storage and transport of high-risk AI systems.
- **Cooperate with authorities.** The same obligations apply to distributors as apply to importers.

6.

Limited-Risk AI

We analyse below the transparency requirements that apply to providers and deployers in relation to limited-risk AI systems under the AI Act (Article 50). The Commission will assess the need to amend this list of limited-risk AI systems every four years (Article 112).

Provider Obligations

- **"Hey, I'm a Chatbot."** Providers must ensure that AI systems intended to directly interact with natural persons, such as chatbots, are designed and developed to inform individuals they are interacting with an AI system. This requirement does not apply where this is obvious for reasonably well-informed, observant, and circumspect individuals, taking into account the circumstances and the context of use.

- **AI-Generated Content.** Providers of AI systems, including general-purpose AI systems and AI systems generating synthetic audio, image, video or text content, must ensure that their systems' outputs are marked in a machine-readable format and detectable as artificially generated or manipulated. Such technical solutions must be effective, interoperable, robust and reliable as far as this is technically feasible. There is little clarity about what this means in practice, so the Commission's guidance will most certainly be helpful.

This obligation does not apply to AI systems performing an assistive function for standard editing or those that do not substantially alter the input data provided by deployers, or the semantics thereof. Again, this exception will need to be further refined in the next few months.

Deployer Obligations

- **Deepfakes.** The AI Act defines deepfakes as AI-generated or manipulated images, audio or video content that resemble existing persons, objects, places, or other entities or events and that would falsely appear to a person as authentic or truthful.

  Businesses using deepfakes in the course of a professional activity must disclose that the content has been artificially generated or manipulated. This obligation does not apply where the use is authorised for law enforcement purposes. Where the content forms part of an evidently artistic, creative, satirical, fictional or analogous work or program, the transparency obligations are limited to disclosing the existence of such generated or manipulated content in an appropriate manner that does not hamper the display or enjoyment of the work.

- **Text.** 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 use is authorised for law enforcement purposes or where the text has undergone a process of human review or editorial control, and where a natural or legal person holds editorial responsibility for the publication of the content.

- **Emotion Recognition and Biometric Categorisation.** Deployers of emotion recognition or biometric categorisation systems, which qualify as high-risk AI systems, must inform individuals exposed thereto of the system's operation. This obligation does not apply to AI systems authorised for biometric categorisation and emotion recognition for law enforcement purposes. Importantly, AI systems that infer individuals' emotions in the workplace or educational institutions are prohibited, unless they are intended to be put in place or into the market for medical or safety reasons.

Transparency, Timing and Format

At the latest, the information regarding the limited-risk AI systems discussed above must be provided in a clear and distinguishable manner at the time of the first interaction or exposure. Other European Union or national laws may impose additional transparency obligations.

The Commission's AI Office will encourage and facilitate the drawing up of codes of practice at the EU level to facilitate the effective implementation of the AI Act's obligations regarding the detection and labelling of artificially generated or manipulated content. The Commission is empowered to adopt implementing acts to approve those codes of practice and, if it considers them inadequate, to adopt specifying common rules for implementing the AI Act's obligations.

GDPR

Where personal data is processed, the GDPR transparency requirements apply in addition to the AI Act obligations. This includes, in particular, transparency about the purposes of the data processing.

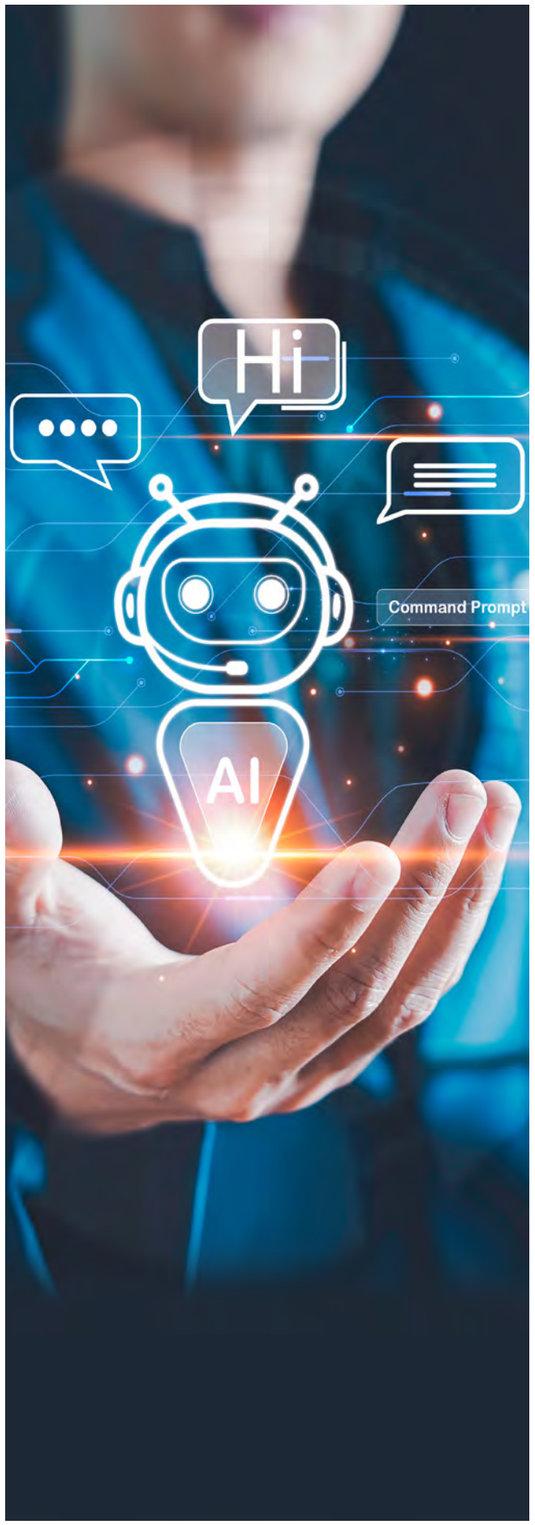

7.

Generative AI

As explained in **chapter 1**, the AI Act generally relies on a risk-based approach. This means that different requirements apply depending on the risk level. GPAI models, however, are a separate category and are subject to specific requirements. These requirements were not part of the Commission proposal in April 2021. They were inserted during the legislative process due to generative AI tools' growing popularity since 2022.

Obligations of Providers of GPAI Models

Providers of GPAI models are required to comply with the following obligations:

- **Technical Documentation for Authorities.** Providers must draft and keep up to date the model's technical documentation, including its training and testing process and the results of its evaluation. Providers must share this information with the Commission's AI Office and the national competent authorities upon request.
  - **General Description.** The technical documentation must include a general description of the GPAI model, including the tasks the model is intended to perform and the type and nature of AI systems in which it can be integrated; acceptable use policies; the release date and distribution methods; the architecture and number of parameters; the modality (e.g., text, image) and format of inputs and outputs; and the license.
  - **Specific Description.** The technical documentation must also include a detailed description of the elements of the GPAI model and relevant information on its development process, including the technical means required for the GPAI model to be integrated in AI systems; the model's design specifications and training process; information on the data used for training, testing and validation; the computation resources used to train the model; and the model's known or estimated energy consumption.
  - **Changes and Specifications.** The Commission may amend and specify the information that needs to be provided in the technical documentation.
- **Documentation for Downstream Providers of AI Systems.** Providers must draft, keep up to date, and supply downstream providers with up-to-date information and documentation on the AI model's capabilities and limitations. Such information must be broadly similar to the information mentioned above. Deployers must take appropriate technical and organisational measures to ensure they use high-risk AI systems that integrate GPAI models in accordance with the downstream provider's instructions for use (see **chapters 4** and **5**).

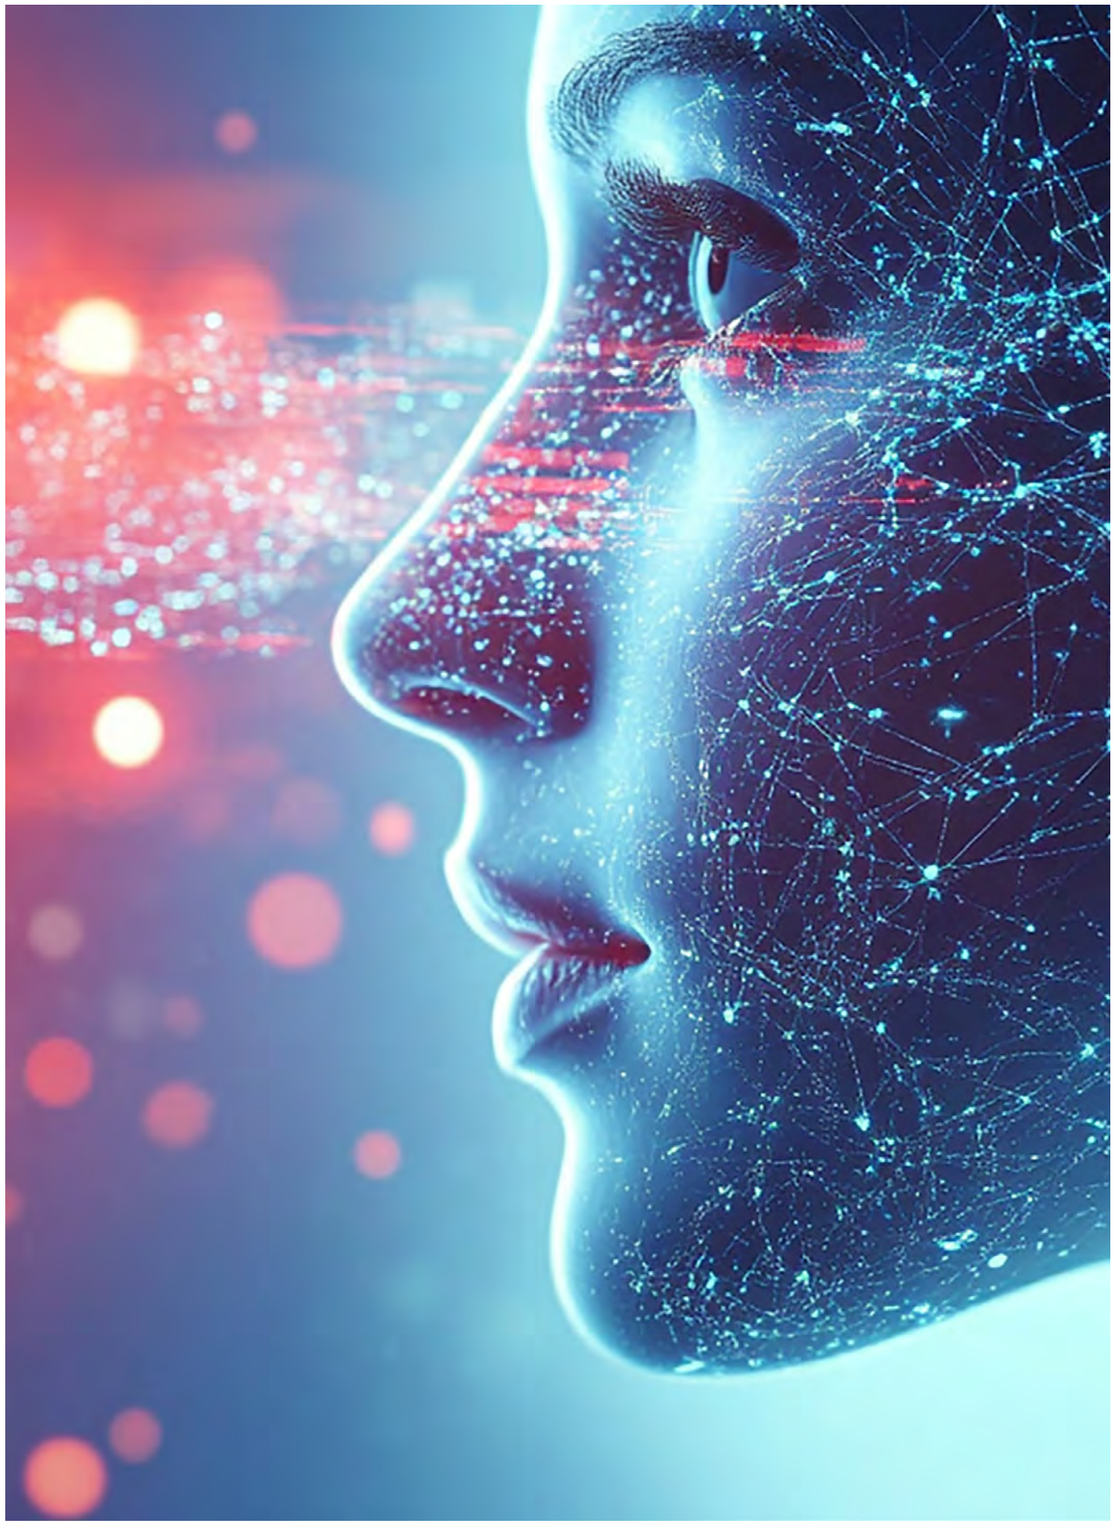

- **Copyright.** Providers must establish a policy to comply with EU law on copyright and related rights, including the EU's Copyright Directive.

- **Information about Content Used for Training Purposes.** Providers must draft and publish a sufficiently detailed summary of the content used for training their AI model, according to a template provided by the AI Office.

- **Cooperation.** Providers must cooperate as necessary with the Commission and the national competent authorities.

- **EU Representative.** A provider must appoint a representative within the EU if it does not have an establishment there.
  - The representative must be appointed by written mandate before placing the GPAI model on the EU market.
  - The representative will manage the technical documentation relevant to its AI model and provide the AI Office and national competent authorities, upon a reasoned request, with all the information and documentation necessary to demonstrate the provider's compliance with its obligations. In addition to, or instead of the provider, the representative can also be addressed by the AI Office or the national competent authorities on all issues related to ensuring AI Act compliance.

Obligations of Providers of Free and Open-License GPAI Models

Providers of free and open-license AI models only have to comply with the copyright and training requirements mentioned above. This exception does not apply if the AI model bears a systemic risk (see below).

Obligations for Providers of GPAI Models with Systemic Risk

A GPAI model bears systemic risk if the provider or the Commission determines that it has ***high-impact capabilities,*** which is defined as having a significant impact on the EU market due to the model's reach or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or society as a whole that can be propagated at scale across the value chain.

- **A GPAI model is presumed to have *high-impact capabilities* where the cumulative amount of computation used for training a GPAI model is greater than 10^25 floating point operations per second (FLOPs).** The Commission may amend this threshold and supplement benchmarks and indicators for this threshold to reflect the state of the art. Providers must notify the Commission without delay and in any event within two weeks of discovery. Providers may present arguments that despite reaching the AI Act's threshold, their models do not present systemic risks due to their specific characteristics.

- **The Commission may also consider that a GPAI model has high-impact capabilities based on various criteria,** including the number of the model's parameters; the quality or size of the data set; the amount of computation used for training the model; the model's input and output modalities; the benchmarks and evaluations of capabilities of the model; whether the model has a high impact on the EU internal market due to its reach; and the number of registered end users. The Commission may amend these criteria.

In addition to the obligations mentioned above, providers of GPAI models with systemic risk are subject to the following requirements:

- **Model Evaluations.** Providers must perform model evaluations in accordance with standardised protocols and tools reflecting the state of the art, including conducting and documenting adversarial testing of the model with a view to identifying and mitigating systemic risks.

- **Risk Mitigation.** Providers must assess and mitigate possible systemic risks at the EU level.

- **Incident Reporting.** Providers must track, document and report serious incidents and possible corrective measures to the AI Office and relevant national authorities.

- **Cybersecurity.** Providers must ensure an adequate level of cybersecurity protection for the model and its physical infrastructure.

Codes of Practice

The AI Office will encourage and facilitate the drawing up of codes of practice at the EU level by May 2025. If a code of practice cannot be finalised by August 2025 or is deemed inadequate by the AI Office, the Commission may provide common rules for the implementation of providers' obligations.

- **Drafting.** The AI Office may invite GPAI model providers to participate in the drawing-up of codes of practice. Other relevant stakeholders (e.g., civil society, industry, academia) may support the process.

- **Monitoring.** The AI Office will ensure that participants in the codes of practice report regularly to the AI Office on the implementation of their commitments and the measures taken and their outcomes. The AI Office and the EU AI Board—the umbrella body that brings together, among others, the national competent authorities and the AI Office—will also regularly monitor and evaluate progress toward the objectives of the codes of practice. The AI Office may invite all GPAI model providers to adhere to the codes of practice.

- **Tools for Compliance.** Until a harmonised standard is published, GPAI model providers with systemic risk may rely on codes of practice to demonstrate compliance with their obligations. Compliance with European harmonised standards grants providers the presumption of conformity. Providers that do not adhere to an approved code of practice or do not comply with a European harmonised standard will need to demonstrate alternative adequate means of compliance for assessment by the Commission.

8.

Innovation and Regulatory Sandboxes

The AI Act contains measures in support of innovation. These measures will be particularly relevant for companies engaged in research and development.

Regulatory Sandboxes

A regulatory sandbox is a tool that allows businesses to explore and experiment with new and innovative products, services, or businesses under a regulator's supervision. It provides innovators with incentives to test their innovations in a controlled environment, allows regulators to better understand the technology, and aims to foster consumer choice in the long run.

Over the past years, the sandbox approach has gained considerable traction across the EU as a means of helping regulators address the development and use of emerging technologies in a wide range of sectors, including fintech, transport, energy, telecommunications and health.

Regulatory Sandboxes in the AI Act (Chapter VI)

The AI Act requires each EU Member State to have at least one operational regulatory AI sandbox (or joint sandboxes with other EU Member States) by August 2, 2026 (see **chapter 2** about the road to full applicability of the AI Act).

Sandboxes should provide a controlled environment for innovation, supporting the development, training, testing, and validation of AI systems under regulatory supervision for a limited period before their placement on the market or entry into service. This should be done according to a sandbox plan agreed on by the prospective providers and the competent authority. Sandboxes may also include testing under real-world conditions within the sandbox environment.

- **Authorities' Role.** The competent authorities must provide guidance, supervision and support within the AI regulatory sandbox to identify risks. They must provide written proof of the activities successfully carried out in the sandbox and an exit report detailing the activities carried out in the sandbox and the related results and learning outcomes. Providers may use such documentation to demonstrate their compliance with the AI Act as part of the conformity assessment process or relevant market surveillance activities. If appropriate, the competent data protection authorities must be associated with the operation of the sandbox.

- **Risk Mitigation.** Any significant risks to health, safety and fundamental rights identified during the AI system's development and testing must be adequately mitigated. The national competent authorities can temporarily or permanently suspend the testing process or participation in the sandbox if no effective mitigation is possible.

- **Liability.** Providers and prospective providers participating in the AI regulatory sandbox remain liable for any damage inflicted on third parties as a result of the experimentation taking place in the sandbox.

- **No Fines.** During this process, no administrative fine should be imposed for infringements of the AI Act so long as the prospective providers observe the sandbox plan and the terms and conditions for their participation and follow in good faith the guidance given by the national competent authority. The same applies regarding infringements of other laws, provided the authorities responsible for such laws are involved in the supervision of the AI system in the sandbox and have provided compliance guidance.

- **Implementing Act.** The Commission will adopt implementing acts to specify detailed arrangements for creating, developing, implementing, operating, and supervising AI regulatory sandboxes. These implementing acts will establish terms and conditions applicable to participants; common principles on the eligibility and selection criteria for participation; and procedures for application, participation, monitoring, exiting and termination.

AI Regulatory Sandboxes and Personal Data

As a general principle, EU data protection law, including the GDPR, remains unaffected by the AI Act's provisions and will also apply to AI regulatory sandboxes. As an exception to that principle, Article 59 of the AI Act provides that personal data lawfully collected for other purposes may be processed in an AI regulatory sandbox solely for the purpose of developing, training and testing certain AI systems in the sandbox.

This approach is very restrictive. It only applies when 10 cumulative conditions are met, including:

- AI systems must be developed to safeguard substantial public interests in areas such as public safety and health, environmental protection, energy sustainability, transport systems and mobility, critical infrastructures and networks, and public services.

- Personal data processed must be necessary for complying with the requirements for high-risk AI systems where those requirements cannot effectively be fulfilled by processing anonymised, synthetic or other non-personal data.

- Any processing of personal data in the sandbox may not affect the data subjects or their data protection rights.

Testing of High-Risk AI Systems in Real-World Conditions Outside AI Regulatory Sandboxes

Providers or prospective providers of specific high-risk AI systems listed in Annex III of the AI Act may test such systems in real-world conditions, outside AI regulatory sandboxes, subject to the conditions outlined below (Article 60). This applies to AI systems dealing with biometrics, critical infrastructures, education and vocational training, employment, worker management, and access to self-employment.

- **Conditions.** The Commission will specify the detailed elements of the real-world testing plan in implementing acts. Testing in real-world conditions can only take place where all the following conditions are met:
  - The (prospective) provider has drawn up a real-world testing plan and submitted it to the market surveillance authority where the testing is to be conducted;
  - The competent market surveillance authority has approved the testing. Such approval may be considered granted in the absence of any response within 30 days, unless otherwise specified by national law;
  - The (prospective) provider has registered the testing in the non-public part of the EU database maintained by the Commission. This does not apply to critical infrastructures;
  - The (prospective) provider conducting the testing is established in the EU or has appointed a legal representative in the EU;
  - The data collected and processed for the purpose of the testing is not transferred to third countries unless appropriate and applicable safeguards under EU law are implemented;
  - The testing lasts no longer than necessary to achieve its objectives and, in any case, no longer than six months (although this may be extended for an additional period of six months);
  - Subjects of the testing who are vulnerable persons due to their age or physical or mental disability are appropriately protected;
  - When (prospective) providers and deployers collaborate on testing, deployers must be given relevant information about all aspects of testing and instructions for use. The (prospective) provider and deployer must agree on their roles and responsibilities to meet testing requirements;
  - The subjects of the testing have given their free and informed consent. Among other things, they must be given information about their rights and the nature and objectives of the testing; any possible inconvenience the testing may cause; and the conditions under which the testing will be conducted;
  - The testing is overseen by the (prospective) providers and deployers through persons who are suitably qualified and have the necessary capacity, training and authority to perform their tasks; and
  - The AI system's predictions, recommendations or decisions can be effectively reversed and disregarded.

- **Withdrawal.** Any subjects of the testing may, without any resulting detriment and without having to provide any justification, withdraw from the testing at any time by revoking their informed consent and may request the immediate and permanent deletion of their personal data. The withdrawal of the informed consent does not affect the lawfulness or validity of activities already carried out.

- **Authorities' Checks.** Market surveillance authorities can require (prospective) providers to supply information, carry out unannounced remote or on-site inspections, and perform checks on the development of the testing and related products to ensure the safe development of testing.

- **Incident Reporting.** (Prospective) providers must report to the competent national market surveillance authority any serious incident identified during the testing and adopt immediate mitigation measures. Failing that, they must suspend the testing until such mitigation takes place or otherwise terminate it. The (prospective) provider must have a procedure for the prompt recall of the AI system upon such termination of the testing. Authorities must be notified accordingly.

- **Liability.** The (prospective) provider must be liable for any damage caused during the testing.

Measures in Support of Small and Medium-Sized Enterprises and Start-Ups

The AI Act requires EU Member States to adopt four key measures to support Small and Medium-Sized Enterprises (SMEs) and start-ups:

- Provide those who have a registered office or branch in the EU with priority access to the regulatory sandboxes;

- Organise specific awareness-raising and training activities on the application of the AI Act tailored to the SMEs' needs;

- Use dedicated channels for communication with them to provide advice and respond to queries about the implementation of the AI Act; and

- Facilitate SMEs' participation in the standardisation development process.

Derogations for Specific Operators

Companies that employ fewer than 10 people and whose annual turnover does not exceed €2 million may comply with certain elements of the quality management system in a simplified manner, subject to additional requirements regarding their size (see **chapter 4** for more detail on quality management systems). The Commission will develop guidelines on this.

9.

Standards, Specifications and Certificates

Standardisation is expected to play a key role in providing technical solutions to ensure compliance with the AI Act given the complexity of the Act's requirements and the technology involved. Many stakeholders are, therefore, closely following and sometimes involved in the development of standards in the field. This is, however, no easy task, and it will require significant efforts to have standards ready for use by August 2026, when most of the AI Act's provisions will come into effect (see **chapter 2** on critical milestones on the road to full applicability of the AI Act for more detail). The European Committee for Standardisation and the European Electrotechnical Committee for Standardisation are working to make the standards available by the end of 2025.

Standards and Specifications

— **Harmonised Standards.** AI systems classified as high-risk and complying with harmonised standards published in the Official Journal of the European Union will be presumed to meet the AI Act's requirements. Such standards will also cover the Act's general transparency requirements under Article 50 (see **chapter 6** for more details on these requirements). To that end, the Commission has issued standardisation requests to European standardisation organisations.

— **Common Specifications.** The Commission is empowered to adopt common specifications for high-risk requirements and limited-risk transparency requirements (Article 41). Harmonised standards, however, take priority. Thus, the Commission may adopt common specifications only if no such standards have been adopted yet; the Commission requested the adoption of such standards but the request has been rejected; or the standards are not delivered in time, are insufficient, or are not compliant with the request.

- High-risk AI systems that are in conformity with common specifications are presumed to be in conformity with the AI Act's requirements covered by those specifications.
- If they are not in conformity, the providers of such systems must show that they have adopted technical solutions that meet a level at least equivalent to those specifications.

Certificates

- **Conformity Assessment Procedure.** Providers must ensure that high-risk AI systems undergo a conformity assessment procedure before placing them on the European market or putting them into service in the EU. Exceptions to this requirement only apply under very strict conditions subject to national market surveillance authorities' and the Commission's review for a limited period of time and only for exceptional reasons of public security or the protection of life, health, the environment, or key industrial and infrastructural assets (Article 46).

  The conformity assessment procedure to be used varies depending on the type of high-risk system and can consist of an internal control or an external control by a Notified Body (NB) under Article 43 (see **chapter 4** for more detail).

- **Certificate Issuance and CE Marking.** If an NB carries out the conformity procedure and determines that the high-risk AI system in question complies with the AI Act's requirements (see **chapter 4** for more detail), the NB will issue an EU technical documentation assessment certificate. The AI Act provides details regarding the required content of such certificates (Annex 7) and their period of validity (Article 44). High-risk AI systems should bear the CE marking to indicate their conformity with the AI Act.

- **Sanctions in Relation to Certificates.** NBs can refuse to issue a certificate if they consider the system non-compliant. They can also suspend or withdraw the certificate or impose restrictions if the system no longer meets AI Act requirements. Providers can avoid such decisions by taking appropriate corrective action within the deadline set by the NB. Providers should also be able to appeal NBs' decisions.

Notifying Authorities

The AI Act requires EU Member States to designate or establish at least one Notifying Authority (NA) responsible for setting up and carrying out the necessary procedures for the assessment, designation, and notification of conformity assessment bodies (CABs) and for their monitoring. EU Member States may decide that such assessment and monitoring may be carried out by national accreditation bodies.

Notified Bodies

- **Role.** NBs must verify the conformity of high-risk AI systems in accordance with the conformity assessment procedure described below (Article 43).

- **From CABs to NBs.** NBs are notified CABs, i.e., bodies that perform third-party conformity assessment activities, including testing, certification, and inspection. To qualify as NBs, CABs must submit a detailed application to the NA of the EU Member State in which they are established. The NA must in turn notify the Commission and the other Member States. CABs may perform the activities of an NB only where neither the Commission nor other EU Member States have raised objections within two weeks to two months of the NA's notification, depending on the documentary evidence submitted by the CAB. If objections are raised, the Commission must enter into consultations with the relevant EU Member States and the CAB and decide whether the CAB can qualify as an NB.

- **Conditions.** NAs may only notify CABs that satisfy the AI Act's requirements (Article 31). The primary goal of these requirements is to ensure that CABs are equipped to conduct independent, objective, impartial, and confidential assessments of high-risk AI systems.

- **CABs established under the law of a non-EU country** may only carry out NBs' activities if they meet the requirements listed in Article 31 or ensure an equivalent level of compliance and the EU has concluded an agreement with the country in question.

10.

Supervision and Enforcement

This chapter introduces the European and national authorities and other relevant actors involved in the supervision and enforcement of the AI Act and provides a brief overview of possible penalties under the AI Act.

Authorities and Other Relevant Actors at the EU Level

- **AI Office.** At the EU level, the AI Act creates the so-called AI Office within the Commission. The AI Office will have the following tasks:
  - **Enforcement.** The AI Office will be responsible for enforcing the AI Act's provisions for providers of GPAI models. However, national competent authorities remain competent vis-à-vis providers and deployers (see below).
  - **Development of Compliance Tools.** The AI Office is generally responsible for developing compliance tools. These include model terms for contracts between providers of high-risk AI systems and third parties providing elements used for or integrated into those systems; templates for deployers' fundamental right to impact assessments of high-risk AI systems; and providers' summaries of the content used for training of general-purpose AI models. The AI Office should also:
    - encourage and facilitate the drawing up of codes of practice at the EU level to contribute to the proper application of the AI Act's requirements regarding general-purpose AI models with systemic risk; and
    - facilitate the effective implementation of the obligations regarding the detection and labelling of artificially generated or manipulated content.
  - **Information About General-Purpose AI Models.** The AI Office may require providers of general-purpose AI models to provide the AI Office with all the information and documentation necessary to demonstrate compliance with the AI Act. For example, the AI Office may ask providers of general-purpose AI models to provide the technical documentation for the model or their authorised representative's mandate.
  - **Reporting Obligations of Providers of General-Purpose AI Models with Systemic Risk.** Providers of such models must report without undue delay to the AI Office relevant information about serious incidents and possible corrective measures.

- **Regulatory Sandboxes.** National competent authorities must inform the AI Office of the establishment and progression of a regulatory sandbox and may ask for the AI Office's support and guidance. Specifically, national competent authorities must inform the AI Office if they suspend the testing process or participation in a regulatory sandbox. The AI Office must make publicly available a list of planned and existing sandboxes and keep it up to date to encourage cross-border cooperation and more interaction in the AI regulatory sandboxes (see **chapter 8** for more detail on regulatory sandboxes).
- **Support for SMEs and Start-ups.** The AI Office must provide standardised templates for areas covered by the AI Act to help SMEs and start-ups comply with the regulations.

— **European AI Board.** The AI Act creates a European AI Board (**Board**) to advise the Commission and EU Member States on the consistent application of the AI Act.

- **Composition.** The Board is composed of one representative per EU Member State. The AI Office and the European Data Protection Supervisor, which enforces the AI Act vis-à-vis the EU institutions, attend the Board's meetings without taking part in the votes. Other authorities, bodies or experts may be invited to the Board's meetings on a case-by-case basis.
- **Tasks.** The Board's primary mission is to advise and assist the Commission and EU Member States to facilitate the consistent and effective application of the AI Act. To that end, the Board has various tasks, such as contributing to the coordination between national competent authorities and harmonisation of national administrative practices; facilitating a common understanding of the AI Act's concepts; assisting national competent authorities and the Commission in developing the organisational and technical expertise required for the Act's implementation; and issuing recommendations and opinions on any relevant matters related to the Act's implementation and its consistent and effective application.

— **Advisory Forum.** An advisory forum will provide technical expertise and advise the Board and the Commission.

- **Composition.** The membership of the advisory forum must represent a balanced selection of stakeholders, including representatives from industry, start-ups, SMEs, civil society, and academia. Membership must also be balanced between commercial interests (including SMEs and larger companies) and noncommercial interests. Several European agencies will be permanent members of the forum. Experts and other stakeholders may be invited.
- **Tasks.** The advisory forum will meet at least twice a year and prepare opinions, recommendations, and contributions upon request of the Board or the Commission. The forum will publish annual reports of its activities.

— **Scientific Panel of Independent Experts.** A scientific panel of independent experts will support the AI Office.

- **Composition.** The panel will consist of experts selected by the Commission based on up-to-date scientific or technical expertise in AI necessary for the panel's tasks. Experts must be independent from providers of AI systems or general-purpose AI models and systems. The Commission, in coordination with the Board, will determine the number of experts and ensure fair gender and geographical representation.
- **Tasks.** The scientific panel will advise and support the AI Office in its tasks. To that end, the panel will have various tasks, such as alerting the AI Office to possible systemic EU-level risks of general-purpose AI models and contributing to the development of tools and methodologies for evaluating capabilities of general-purpose AI models and systems. In addition, EU Member States may ask the panel to support their enforcement activities.

Authorities and Other Relevant Actors at the National Level

— **Market Surveillance Authorities.** Each EU Member State must establish or designate at least one market surveillance authority responsible for ensuring compliance with the AI Act. These authorities will be primarily responsible for enforcing the AI Act, including ensuring compliance of downstream GPAI model providers

and deployers with the transparency requirements detailed in **chapter 6**.

- **Notifying Authorities.** Each EU Member State must establish or designate at least one notifying authority responsible for designating and monitoring conformity assessment bodies. Such bodies may become notified bodies responsible for the performance of third-party conformity assessment activities, including testing, certification, and inspection (see **chapter 9** on standardisation in the AI Act for more detail).
- **Guidance and Advice.** Market surveillance authorities may provide guidance and advice on the implementation of the AI Act, in particular to SMEs, including start-ups. In so doing, national competent authorities must take into account the guidance and advice of the Board, the Commission or any other relevant authority.
- **Sufficient Resources.** EU Member States must ensure that market surveillance authorities are provided with adequate technical, financial and human resources, and with infrastructure to fulfil their tasks. EU Member States will need to report to the Commission on the status of their authorities' resources mid-2025 and once every two years thereafter. The Commission will share this report with the Board for discussion and possible recommendations.

Penalties

Under the AI Act, EU Member States must set penalties for infringements of the AI Act and take all measures necessary to ensure that they are implemented.

Penalties must be effective, proportionate, and dissuasive and must take into account:

- the nature, gravity and duration of the infringement and its consequences;
- whether other market surveillance authorities have already fined the operator for the same infringement; and
- the size and market share of the operator committing the infringement.

Administrative fines will be based on an undertaking's global revenue if the amount exceeds the fine cap. Fines include:

- Up to €35 million or up to 7% of revenue for noncompliance with the AI Act's requirements regarding prohibited AI systems;
- Up to €15 million or up to 3% of revenue for noncompliance with the Act's requirements regarding limited and high-risk AI and general-purpose AI models; and
- Up to €7.5 million or 1% of revenue for supplying incorrect, incomplete, or misleading information when addressing authorities' requests.

Contact our Teams

For any additional information on AI or data-related issues under EU law, please contact our teams in Brussels, Frankfurt and London.

BRUSSELS

Anne Vallery

Partner-in-Charge
anne.vallery@wilmerhale.com

Frederic Louis

Partner
frederic.louis@wilmerhale.com

Itsiq Benizri

Counsel
itsiq.benizri@wilmerhale.com

FRANKFURT

Dr. Martin Braun

Partner
martin.braun@wilmerhale.com

Prof. Dr. Hans-Georg Kamann

Partner
hans-georg.kamann@wilmerhale.com

LONDON

Cormac O'Daly

Partner
cormac.o'daly@wilmerhale.com

Connect with us

wilmerhale.com