

EU AI Act Handbook

June 2025

Contents

Chapter 01	Introduction	2
Chapter 02	Scope and application of the EU AI Act	6
Chapter 03	Core definitions	10
Chapter 04	AI literacy	14
Chapter 05	Prohibited AI practices	17
Chapter 06	Classification of AI systems as “high-risk”	21
Chapter 07	Requirements for high-risk AI systems	26
Chapter 08	Obligations relating to high-risk AI systems	31
Chapter 09	Notifying authorities and conformity assessment bodies	37
Chapter 10	Standards, conformity assessments, certificates, and registration	41
Chapter 11	Transparency obligations for certain AI systems	47
Chapter 12	GPAI models – Classification rules	53
Chapter 13	GPAI models – General obligations of providers of GPAI models	58
Chapter 14	GPAI models – General obligations of providers of GPAI models with systemic risk	62
Chapter 15	GPAI models – Codes of practice	65
Chapter 16	Measures designed to support innovation	69
Chapter 17	Regulatory framework	75
Chapter 18	EU database for high-risk AI systems	78
Chapter 19	Monitoring and oversight	83
Chapter 20	Codes of conduct and guidelines	88
Chapter 21	Implementing regulations	91
Chapter 22	Penalties	94
Chapter 23	Commencement and timing	97
Chapter 24	Overlap between the EU AI Act and other relevant laws and guidelines	101
Glossary		105
Contributors		113
About White & Case		115

Chapter 01

Introduction

Executive summary

The EU AI Act is a complex law that is, in places, hard to understand. This Handbook is designed to help businesses navigate these challenges by providing a pragmatic analysis of the issues they are likely to face under the EU AI Act. It emphasises practical compliance over legal theory, offering actionable guidance and insights wherever possible.

Analysis

Over recent years, AI has moved from experimental corners of research labs into the operational core of businesses across the EU and beyond. From supply chain optimisation to personalised marketing, AI is reshaping industries at breakneck speed. Yet, as innovation races ahead, regulation is just getting started – and nowhere is this catch-up effort more evident than in the EU AI Act.

The EU AI Act is wide-ranging (applying in a variety of different contexts), top-down (with many of its requirements being imposed directly onto businesses), and crucially, still in flux (with many of its requirements and even some of its core definitions being subject to ongoing review and amendment). While its intent is commendable – to ensure trustworthy, safe, and rights-respecting AI – its implementation raises a practical conundrum for businesses: **How do you approach a law that is vague in key areas, is still evolving, and inconsistently understood even among experts?**

This Handbook is designed to help businesses answer that question. It is not a legal commentary, a treatise, or a speculative essay on the future of AI governance. Instead, it is unapologetically pragmatic. This Handbook is designed for businesses who need to understand the EU AI Act, and need clarity, structure, and direction. We accept that there are substantial areas in which nobody can claim to know with certainty how the EU AI Act will be applied by

courts and regulators. Nevertheless, wherever possible, this Handbook seeks to give guidance on likely positions, often reasoning by analogy based on White & Case’s deep experience of interpreting vague and uncertain provisions in other EU laws.

Where the EU AI Act is ambiguous, we aim to be clear. Where it is high-level, we aim to be grounded. And where it is theoretical, we focus on the interpretations that the EU’s courts are most likely to adopt. We do not claim to solve all uncertainties – nobody can, not yet – but we provide a map that helps you move forward today, not just prepare for tomorrow.

Unclear terms, undefined risks

If you’ve tried to read the EU AI Act, you know that it can be pretty dense in places. Spanning hundreds of pages, laden with legal terms, cross-references, and numerous mentions of other EU laws, the EU AI Act regulates two separate but distinct concepts: AI systems, and GPAI models.

- **AI systems** – The EU AI Act categorises AI systems into risk tiers (minimal, limited, high, and prohibited) and lays out different requirements accordingly. This sounds structured in theory, but in practice there are gaps, grey areas, and moving goalposts. It references systems that use techniques such as machine learning, logic-based approaches, and statistical methods. But these concepts are so broad that they encompass everything from advanced neural networks to simple rule-based decision trees – arguably even something as simple as auto-correct or predictive text – blurring the line between what *is* and *is not* AI under the law.

For instance, if a company is developing an automated email routing tool based on a series of if/then rules designed to react to the content of each email and route it accordingly, how complex do those rules have to be before the tool is an “AI system” within the meaning of the EU AI Act? A very simple email routing tool based solely on key words in the email subject is likely not an “AI system”, while a very complex email routing tool that uses language analysis to determine from context whether an email is a complaint, an order renewal, a request to renegotiate, etc., and then routes

emails to the relevant department is more likely to be an “AI system”. Even though the Commission sought to clarify some of the uncertainties on this issue in its [Guidelines on the definition of AI systems](#), the boundary between such examples remains difficult to delineate. For this reason, this Handbook provides detailed commentary on the core elements a business will need to consider in order to help determine whether a given tool is (or is not) an “AI system” (see Chapter 3).

- **GPAI models** – These are models that are capable of performing a wide range of tasks in lots of different contexts (e.g., large language models or foundation models trained on vast datasets). But as with AI systems, there are questions about what exactly constitutes a GPAI model. For instance, Rec.97 states that one GPAI model can be “fine-tuned” into a new GPAI model. Exactly how much fine-tuning needs to happen before this change occurs? And how can a business realistically track whether such a change has taken place? Unfortunately, the answers to many of these questions remain unclear. For businesses using or building on GPAI models, this lack of clarity poses significant compliance uncertainty.

Worse still, many core concepts – like “substantial modification” of an AI system, “reasonably foreseeable misuse”, or “sufficiently representative datasets” – are either undefined, or are defined in a manner that creates as many questions as it answers (e.g., the definition of “reasonably foreseeable misuse” relies on an intended purpose but does not explain whether that intention is determined by the provider, the deployer, and/or a third party. It also does not explain whether that intent needs to be stated in advance or whether oblique intent can be inferred). Unfortunately, businesses are left guessing. And these ambiguities are not just academic; they have real operational consequences for businesses trying to develop and/or use AI technologies in a compliant manner.

The need for practical guidance

Businesses need practical guidance because the stakes are high. The penalties for non-compliance with the EU AI Act are significant – up to a maximum of €35 million or 7% of global annual turnover, whichever is greater. To put it another way, these figures are 75% higher than the

already eye-watering maximum penalties in the GDPR. In addition to the risk of financial exposure, the reputational risk of being branded as non-compliant is also significant, especially in a market where trust and transparency are quickly becoming competitive differentiators.

Yet, most businesses do not have the time, expertise, or internal legal capacity to spend hours researching the answers to questions such as “How do I know if my AI system is ‘high-risk’?”; “What transparency notices do I need to start preparing?”; “Do I need to retrain my teams on AI literacy?” and so on. For that reason, we created this Handbook to provide businesses with practical insights into the EU AI Act, and guidance on how to approach compliance.

Embracing the grey areas

To be clear: **This Handbook does not pretend to offer final answers where none yet exist.**

In some areas – including interactions with regulators, interpretations of risk categories, or evolving technical standards – the EU AI Act is either silent, unclear, or defers to future harmonised standards, which are themselves still being developed. These grey areas can make the task of implementing an EU AI Act compliant project seem overwhelming.

But we are also not powerless. This is not the first time that we have all faced uncertain or unclear provisions of EU law (and sadly, it will not be the last). But decades of case law and regulatory guidance have given us a range of tools that we can use to anticipate the ways in which the EU AI Act will be interpreted and applied. In particular:

- **Acknowledge uncertainty where it exists** – As a starting point, it is essential for lawyers and other advisors to be honest about the areas in which the law simply is not clear. This will help business teams to better understand that many of the compliance decisions that need to be taken are essentially an exercise in balancing risk.
- **Accept that interpretations will change over time** – As with many other EU laws, it is inevitable that the interpretations applied by courts and regulators to the EU AI Act will evolve with time. For example, we have

seen this evolution under the GDPR – transparency measures and data transfer mechanisms that were considered lawful when the GDPR came into effect in May 2018 have (as a result of case law and changes in the positions adopted by regulators) become non-compliant, and businesses have needed to adapt. Similar changes over time are likely to happen under the EU AI Act.

- **Be aware of parallels with other EU laws** – As set out in detail in Chapter 24, the EU AI Act does not exist in a vacuum – it operates in concert with a host of other EU laws, many of which have been through similar exercises of clarifying uncertainties, and these can help us draw parallels with the EU AI Act. In particular (as discussed in Chapter 3), the CJEU has adopted a purposive interpretation of unclear provisions of other laws (such as the GDPR), and we can apply the same approach to interpreting some of the unclear provisions of the EU AI Act.
- **Seek progress over perfection** – Waiting for 100% clarity before taking action is unlikely to be a successful strategy. Businesses should start building their compliance foundations as early as possible, knowing they can adapt. In the event that a business faces a regulatory investigation under the EU AI Act, it is infinitely preferable to demonstrate a compliance program in progress than to have nothing to show. This Handbook is designed to help businesses get started as quickly and easily as possible.
- **Adopt defensible positions** – Inevitably, each business will face risk calls on how best to balance compliance obligations under the EU AI Act with the need to be able to develop and/or use AI to achieve its business goals. For instance, each business needs to decide which of its technologies are AI systems, and which are not. Given how vaguely the EU AI Act defines “AI systems”

(see Chapter 3), businesses will need to adopt internal definitions that strike a balance between being too broad (risking treating certain technologies as being in-scope when they might not be) and too narrow (risking non-compliance). When striking such a balance, it is essential for businesses to be able to justify and defend the choices they have made, and to document those choices. By doing so, businesses can demonstrate to any regulator the reasonableness of the choices that they have made, especially when those choices had to be made before thorough guidance was available.

- **Keep up-to-date** – Unfortunately, in this rapidly changing regulatory environment, there is no substitute for keeping up-to-date with regulatory trends, legal guidance, and real-world implementation feedback. Regularly updated resources such as this Handbook (and our [AI Watch global regulatory tracker](#) for those looking beyond the EU) are designed to make this process easier.
- **Enforcement will take some time** – Immediate enforcement on day one is highly unlikely for most businesses. In addition, very large penalties (penalties of multiple millions of Euros, or penalties based on a percentage of turnover) are unlikely for a first offence. As we have seen with other EU laws (especially those that rely on an element of enforcement by national regulators), it takes time for regulators of new laws to be set up and to get to grips with their new powers, and where the correct interpretation is unclear they may be hesitant to issue very large penalties, for fear of being overturned on appeal (whereas businesses might not risk appealing smaller penalties). In addition, as noted in our EU AI Act enforcement timeline, the start of enforcement of the EU AI Act is staggered over several years. As a result, it seems likely that enforcement will take some time to ramp up.

What this Handbook is (and isn't)

Theory without implementation is meaningless. The core goal of this Handbook is to translate legislative language into principles, guidance, and commentary that can be more easily digested and applied in practice.

This Handbook *is*:

- A business-focused interpretation of the EU AI Act.
- A source of commentary on the ways in which the EU AI Act is likely to be applied in practice, and the areas in which it remains unclear.
- A living document that will evolve with the law and its interpretation and make it easier to keep up-to-date with the EU AI Act.

This Handbook *is not*:

- A substitute for legal advice.
- An exhaustive treatise on the EU AI Act, or the philosophical underpinnings of AI regulation in general.
- A static text – it will need to evolve as regulatory guidance, case law, and technical standards emerge.

Where to go from here

Businesses in almost all sectors need to ensure that they are aware of the developments under the EU AI Act (including the appointment of national regulators, the publication of new guidance, and the emergence of case law, all of which will affect the impact that the EU AI Act has on businesses). It is essential for businesses to keep up-to-date with these developments, in order to identify new opportunities and new potential business risks. To that end, this Handbook will be regularly updated, to help businesses keep abreast of developments in this rapidly changing regulatory environment.

Key contacts



Tim Hickman

Partner, London

E tim.hickman@whitecase.com



Sylvia Lorenz

Partner, Berlin

E sylvia.lorenz@whitecase.com



Jenna Rennie

Partner, London

E jenna.rennie@whitecase.com



Clara Hainsdorf

Partner, Paris

E chainsdorf@whitecase.com

Chapter 02

Scope and application of the EU AI Act

Executive summary

The EU AI Act applies to a wide range of roles across the AI value chain, such as providers (who develop AI systems or GPAI models) and deployers (who use AI systems or GPAI models).

Territorially, the EU AI Act applies both within the EEA and outside the EEA in a very wide range of circumstances. Businesses are at risk of being subjected to the EU AI Act, even if they are not intending to do business in the EEA.

The conceptual scope of the EU AI Act is limited to areas within the EU's legislative competence and is subject to a number of limitations and exemptions.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.2** of the EU AI Act – specifically, the in-scope roles, territorial scope, and conceptual scope of the EU AI Act. This Chapter also includes insights on the relevant definitions in **Art.3**, to the extent that those definitions relate to the scope of the EU AI Act.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **AI system** (Rec.12; Art.3(1)) – A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.
- **GPAI model** (Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market, and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.
- **Deployer** (Rec.13; Art.3(4)) – Any organisation using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.
- **Placing on the market** (Art.3(9)) – The first making available of an AI system or a GPAI model on the EEA market.
- **Putting into service** (Art.3(11)) – The supply of an AI system for first use directly to the deployer, or for the provider's own use in the EEA, for its intended purpose.

A full list of defined terms can be found in the [Glossary](#).

Analysis

In-scope roles – The EU AI Act applies to actors in a wide range of different roles in the AI value chain. Unlike, for example, the GDPR (which essentially regulates only two roles – controller and processor), the EU AI Act regulates six distinct roles. The roles to which the EU AI Act applies are as follows:

- **Provider** (Recs.21 and 22; Arts.2(1)(a), (c), and 3(3)) – Any organisation that develops an AI system/GPAI model, or that has an AI system/GPAI model developed and places it on the market or puts the AI system into service under its own name or trademark. Providers do not have to be established or located in the EEA, nor do they have to necessarily place an AI system on the EEA market (provided the output of the AI system is used in the EEA).

- **Deployer** (Recs.21 and 22; Arts.2(1)(b), (c), and 3(4)) – Any organisation that uses an AI system under its own authority, except where the AI system is used in the course of a personal non-professional activity. Deployers do not have to be established or located in the EEA, nor do they have to necessarily place an AI system on the EEA market (provided the output of the AI system is used in the EEA).
- **Importer** (Recs.83 and 84; Arts.2(1)(d) and 3(6)) – Any organisation located or established in the EEA that places an AI system on the market which bears the name or trademark of an entity established outside of the EEA is an “*importer*”.
- **Distributor** (Recs.83 and 84; Arts.2(1)(d) and 3(7)) – Any organisation (other than a provider or importer) that provides AI systems/GPAI models for distribution or use on the EEA market is a “*distributor*”. The distributor does not need to be the first organisation in the AI value chain that releases the AI system/GPAI model to the EEA market.
- **Product manufacturer** (Rec.87; Arts.2(1)(e) and 25(3)) – The concept of a “*product manufacturer*” is not explicitly defined in the EU AI Act (instead, it is defined in the EU harmonisation legislation listed in Annex I to the EU AI Act – see Rec.87). Product manufacturers are within the scope of the EU AI Act when they place an AI system on the EEA market together with their own products and under their own name or trademark. In certain cases, a product manufacturer will be deemed to be the provider of a high-risk AI system, where the high-risk AI system is placed on the market together with the product under the name or trademark of the product manufacturer, or is put into service under the name or trademark of the product manufacturer after the product has been placed on the market.
- **Authorised representative** (Rec.82; Arts.2(1)(f) and 3(5)) – Authorised representatives are intermediaries appointed by providers outside of the EEA. An “*authorised representative*” is any organisation in the EEA that has accepted a written mandate from the provider to carry out the provider’s obligations with respect to the EU AI Act.

In addition, the EU AI Act recognises:

- **Affected person** (Recs.20 and 171; Art.2(1)(g)) – The term is not explicitly defined in the EU AI Act but appears to mean individuals affected by AI.
- **Operator** (Rec.22; Art.3(8)) – The term is a catch-all term for providers, product manufacturers, deployers, authorised representatives, importers, and distributors. It is worth noting that Art.2 does not set out separate rules governing the applicability of the EU AI Act to operators. Consequently, when that term is used, the applicability of the EU AI Act appears to depend on whether that operator is a provider, deployer, etc.

Territorial scope – The EU AI Act has an aggressive approach to extraterritoriality. As a result, businesses operating outside the EEA are nevertheless at risk of being subject to the EU AI Act, even if they are not intending to do business in the EEA.

The EU AI Act applies to any organisation established or located in the EEA that uses AI systems. This includes providers, deployers, importers, distributors, and product manufacturers based in, or operating in, the EEA.

In addition, the EU AI Act has a broader approach to extraterritoriality than comparable EU regulatory laws (e.g., the GDPR). Whereas the GDPR focuses on entities outside the EEA that offer services to or monitor individuals in the EEA, the EU AI Act applies to any provider or deployer of AI systems if the output is used in the EEA, seemingly regardless of intent.

Specifically, Rec.22 indicates that the EU AI Act should apply to providers and deployers outside the EEA if the output of their AI systems is intended to be used in the EEA. However, this seems inconsistent with Art.2(1)(c), which states that “[The EU AI Act] applies to: [...] providers and deployers of AI systems that have their place of establishment or who are located in a third country [i.e., outside the EEA], where the output produced by the system is used in the [EEA].” This language removes the element of intent and instead appears to mean that the EU AI Act applies if the output is used in the EEA, regardless of whether this was intended. As a result, businesses operating outside the EEA are nevertheless at risk of being subject to the EU AI Act, even if they are not intending to do business in the EEA.

A further complication stems from the concept of “*affected persons*”. Under the GDPR, if a business does not fall within any of the territorial scope tests set out in Art.3 GDPR, then that business is not subject to the GDPR, even if some of the affected data subjects are located in the EEA. But Art.2(1)(g) of the EU AI Act states that the EU AI Act applies to affected persons located in the EEA. Although the wording is unclear, it is possible that this is intended to mean that even where a business passes none of the other tests for applicability of the EU AI Act, any affected person located in the EEA may still be able to exercise their rights under the EU AI Act against that business.

The EU AI Act is a “*Text with EEA Relevance*”. This means that, under the EFTA Treaty, the non-EU EEA States (i.e., Iceland, Liechtenstein, and Norway) will need to implement national laws to give effect to the EU AI Act. However, in practice the rules will apply in those states in functionally the same way that they apply in the EU.

Limitations on scope – The scope of the EU AI Act is limited to areas that are within the EU’s legislative competence and is subject to a number of limitations and exemptions.

The EU AI Act contains a number of carve-outs and limitations on its applicability. In particular:

- For high-risk AI systems that are covered by EU harmonisation legislation listed in Annex I(B), only certain provisions of the EU AI Act apply (Art.2(2)).
- The EU AI Act does not apply to areas outside the scope of EU law. It also does not affect the competences of EEA states concerning national security. It also does not apply to AI systems used exclusively for military, defence, or national security purposes (Art.2(3)).

- The EU AI Act does not apply to AI systems used by public authorities of countries outside the EEA within the framework of international cooperation or agreements for law enforcement and judicial cooperation, provided that appropriate protections for individuals’ rights are in place (Art.2(4)).
- The EU AI Act does not affect the provisions on liability of intermediaries in the DSA (Arts.2(5)).
- The EU AI Act does not apply to AI systems or AI models that are designed and used solely for scientific research and development (Arts.2(6)).
- The EU AI Act does not affect the application of the GDPR or the e-Privacy Directive (Art.2(7)) or EU laws regarding consumer protection and product safety (Art.2(9)).
- The EU AI Act does not apply to any research, testing, or development of AI systems or AI models prior to their being placed on the market or put into service – but this exclusion does not extend to testing in real-world conditions (Art.2(8)).
- The EU AI Act does not apply to deployers using AI systems exclusively for personal, non-professional activities (Art.2(10)).
- The EU AI Act does not prevent EEA Member States from making their own laws protecting workers from the impact of AI (Art.2(11)).
- The EU AI Act does not apply to AI systems released under free and open-source licences, unless they are high-risk AI systems, prohibited AI systems, or AI systems that are subject to the transparency obligations in Art.50 (Art.2(12)).

Context and illustrations



Commentary: Definition of “provider”

The syntax in the definition of “provider” (Art.3(3)) is unclear. A literal reading of the definition would indicate that a business can only be a provider if it: (i) develops an AI system or AI model (or has one developed); **and** (ii) places that AI system or AI model on the market or puts it into service under its own name or trademark. However, this interpretation appears to leave some gaps. For example, where one party develops an AI system, and a different party places it on the market, neither of them satisfies the definition of “provider” (meaning that such an AI system would have no provider). Alternatively, the “and” above could possibly be read as an “or”, but that would also create difficulties because it would result in multiple parties being the provider of the same AI system, leading to uncertainty about who is responsible for which compliance obligations in relation to that AI system. Businesses are likely to face substantial uncertainty in this area until regulatory guidance or court decisions addressing these issues are published.



Example: Extraterritorial application of the EU AI Act

Company X is an advertising agency based in Japan. Company X has implemented third party AI systems into its workflow and uses those systems to generate branding concepts for clients. Company Y is a customer and is based in Argentina. Company Y engages Company X to create branding for one of its products, and Company X does so, using AI systems to generate some elements of the branding. Company Y loves the branding and pays for the work.

A year later, Company Y opens a new office in Spain, where it uses the branding created by Company X. Under the provisions of the EU AI Act, Company X is a deployer of the AI system that it used to develop the branding. The branding is the “output” of that AI system. Company Y has used that output in the EEA. Therefore, under a literal interpretation of Art.2(1)(c), Company X (which is solely based in Japan, and was not intending to do business in the EEA) is now subject to the EU AI Act. It also does not appear that there is any way that Company X can avoid this fate. Even if Company X contractually prohibited its customers from using its branding in the EEA, it seems that this would have no effect because Art.2(1)(c) does not appear to take intent into account.



Commentary: Overlap between the EU AI Act and other EU laws

The overlap between the EU AI Act and other EU laws is complex. On the one hand, Art.2 is reasonably clear that the EU AI Act is intended to be “without prejudice to”, or otherwise is intended not to affect, overlapping provisions in other EU laws, including the GDPR, the e-Privacy Directive, and the DSA. However, in practice the relationship between those laws and the EU AI Act is likely to require substantial clarification in the form of court decisions and regulatory guidance. In particular, where a business uses an AI system to process personal data, and something goes wrong, is that business potentially exposed to parallel liability under both the GDPR and the EU AI Act, in respect of the same incident? If so, that would give rise to two parallel sets of enforcement proceedings and appeals, and ultimately an aggregate maximum fine of **the greater of €55 million or 11% of worldwide turnover** (being the GDPR maximums of €20 million or 4%, plus the EU AI Act maximums of €35 million or 7%) in addition to any damages payable under either law.

Chapter 03

Core definitions

Executive summary

The EU AI Act contains a series of core definitions that are critical to understanding the EU AI Act's scope and application. Many of these definitions have intentionally been drafted with an element of flexibility to allow them to apply to rapidly evolving AI technologies. However, with flexibility sometimes comes uncertainty, and these definitions are not always clear.

Businesses have little choice but to attempt to apply these definitions to their relevant activities by extrapolating from interpretations adopted in other analogous EU laws, and hope that additional clarity will be provided by upcoming regulatory guidance and case law.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.3** of the EU AI Act, which provides a list of definitions. This Chapter analyses several of the core definitions, considers the ways in which those definitions leave room for doubt or uncertainty, and analyses how they are likely to be interpreted in future case law and regulatory guidance.

Key defined terms

The key defined terms used in this Chapter are assessed in detail in the Analysis section below. A full list of defined terms can be found in the [Glossary](#).

Analysis

"AI system" (Rec.12; Art.3(1)) – The term *"AI system"* is foundational to any understanding of the EU AI Act, as many of the core compliance obligations arise in relation to AI systems. However, the definition in Art.3(1) of the EU AI Act leaves substantial room for uncertainty. The key elements of that definition are as follows:

- **AI system means "a machine-based system..."**
 - Rec.12 indicates that this *"refers to the fact that AI systems run on machines"*, but it is not clear whether this means an AI system must run exclusively on machines. For example, if a system is reliant on some level of interaction from humans for its operation and upkeep, is that system still *"machine-based"*? The Commission issued [guidelines on the definition of AI systems](#) in February 2025 (the "AI Systems Guidelines") which clarify that *"machine-based"* includes a variety of hardware and software systems, but offer no explicit clarification on the status of systems that require some level of human interaction or upkeep. Nevertheless, based on the expansive interpretation applied by the Commission, it seems likely that systems that include some degree of human involvement would still fall within the definition, provided that the other elements (below) are satisfied.
- **"...designed to operate with varying levels of autonomy..."** – According to Rec.12, *"autonomy"* means *"some degree of independence of actions from human involvement and [the ability] to operate without human intervention"*. The AI Systems Guidelines indicate that systems that require *"full manual human involvement"* are out of scope. On a literal reading, it seems that any degree of autonomy would be sufficient to satisfy the definition, although the AI Systems Guidelines set a slightly higher bar, indicating that the minimum requirement is *"some reasonable degree of independence of actions"*. The word *"varying"* presumably indicates that a wide range of levels of autonomy are in-scope (as opposed to meaning that the level of autonomy displayed by a system must vary over time).
- **"...that may exhibit adaptiveness after deployment..."** – Rec.12 clarifies *"adaptiveness"* that *"refers to self-learning capabilities, allowing the system to change while in use"*. On a literal reading, the word *"may"* seems to mean that adaptiveness is not a strict

requirement for a system to be an AI system – i.e., a system might have adaptiveness, or might not, and this would not necessarily prevent it from being an AI system. The AI Systems Guidelines confirm that an AI system “does not necessarily have to possess adaptiveness”.

□ **“...for explicit or implicit objectives...”** – On a literal reading, this simply seems to mean “all objectives”. Rec.12 clarifies that this term captures both systems that are directed to produce outputs that are specified at the outset (which the AI Systems Guidelines describe as “clearly stated goals that are directly encoded by the developer into the system”) and systems that produce outputs that are different from the intended purpose of the system (which the AI Systems Guidelines describe as “goals that are not explicitly stated but may be deduced from the behaviour or underlying assumptions of the system”).

□ **“...infers, from the input it receives, how to generate outputs...”** – Rec.12 emphasises that the ability to “infer” is an essential characteristic of AI systems. The term “infer” indicates that the system must independently derive outputs from the inputs it receives. The AI Systems Guidelines devote substantial attention to the meaning of the term “infer” and give examples of AI systems including “image classification systems trained on a dataset of images... medical device diagnostic systems trained on medical imaging labelled by human experts, and fraud detection systems that are trained on labelled transaction data.” These are contrasted with examples of systems that predict stock prices or temperatures based on historical averages, using “a basic statistical learning rule” – which are not deemed to be AI systems, on the basis that they “have the capacity to infer in a narrow manner but may nevertheless fall outside of the scope of the AI system definition because of their limited capacity to analyse patterns and adjust autonomously their output.” Notwithstanding the detail in the examples, the precise distinction between a system that can “infer” and one that cannot remains elusive.

□ **“...such as predictions, content, recommendations, or decisions...”** – The term “such as” is an indication that these are merely examples of types of output which are not intended to be exhaustive. A system could produce outputs that fall into none of these

categories and still potentially be an AI system. The AI Systems Guidelines state that AI systems “differ from non-AI systems in their ability to generate outputs like predictions, content, recommendation, and decisions in that they can handle complex relationships and patterns in data.”

□ **“...that can influence physical or virtual environments”** – This expression is unclear, and the Recitals do not bring further explanation. At a minimum, it seems that a system must have an identifiable influence of some kind (whether in a physical or a virtual context) to qualify as an AI system. The AI Systems Guidelines state that the reference to physical or virtual environments “indicates that the influence of an AI system may be both to tangible, physical objects (e.g., robot arm) and to virtual environments, including digital spaces, data flows, and software ecosystems.”

“General-purpose AI model” (Recs.97 – 99; Art.3(63)) – The term “general-purpose AI model” or “GPAI model” is essential to any understanding of the EU AI Act, as a substantial number of compliance obligations arise in relation to GPAI models. The key elements of that definition are as follows:

□ **A GPAI model is “an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale...”** – Rec.97 explains that “AI models” are essential components of AI systems, but do not constitute AI systems on their own. AI models are essentially building blocks on which AI systems are built. They require additional features (e.g., a user interface) before they can become AI systems. Rec.99 states that “[l]arge generative AI models are a typical example [of a GPAI] model”.

□ **“...that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market...”** – Rec.97 emphasises that the concept of GPAI models is rooted in the generality of such models, and their ability to competently perform a wide range of distinct tasks. However, the expression “competently perform” is only used on three occasions in the EU AI Act and is not fully explained. If a model has significant generality but poor performance (so that it can perform a wide range of tasks badly – i.e., it cannot do so **competently**), then it is uncertain whether that model is a GPAI model.

□ **“...and that can be integrated into a variety of downstream systems or applications...”** – Rec.101 explains that, because GPAI models “may form the basis for a range of downstream systems”, the providers of those downstream systems need a good understanding of the capabilities of the relevant GPAI models. Therefore, the EU AI Act imposes transparency obligations on the providers of GPAI models.

□ **“...except AI models that are used for research, development or prototyping activities before they are placed on the market”** – This exemption is not further clarified. Rec.97 merely explains that when a GPAI model is placed on the market for any purpose other than research, development, or prototyping, the exemption ceases to apply.

“Provider” (Art.3(3)) – A “provider” is one of the key compliance roles assigned under the EU AI Act in relation to AI systems and GPAI models.

□ **Provider means any organisation “that develops an AI system or a general-purpose AI model, or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge”** – As noted in the commentary in Chapter 2, this definition suggests that a business will only be a provider if it: (i) develops an AI system or AI model (or has one developed); and (ii) places that AI system or AI model on the market or puts it into service under its own name or trademark. If the “and” is read conjunctively, this would mean that a business that does either (i) or (ii) (but not both) cannot be a provider – meaning that in many cases there would be no provider. Alternatively, reading the “and” disjunctively (so that any business that does either (i) or (ii) is a provider) would result in multiple parties being the provider of the same AI system, leading to uncertainty about who is responsible for which compliance obligations in relation to that AI system.

“Deployer” (Rec.13; Art.3(4)) – A “deployer” is one of the key compliance roles assigned under the EU AI Act in relation to AI systems.

□ **Deployer means “a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity”** – On a literal reading, this definition is

comparatively straightforward. Any business using an AI system is a deployer (subject to an exception that seemingly does not apply to any business). But on further analysis, it appears that a business that satisfies the definition of a “provider” in developing an AI system and/or placing it on the market becomes a “deployer” when it then uses that AI system – meaning that such a business will need to ensure that it fulfils the compliance obligations of both a provider and a deployer in relation to that AI system.

“Biometric identification” (Rec.15; Art.3(35)) – The EU AI Act places significant emphasis on regulating biometric identification systems due to their potential impact on privacy and security.

□ **Biometric identification means “the automated recognition of physical, physiological, behavioural, or psychological human features [...] by comparing biometric data of that individual to biometric data of individuals stored in a database”** – As per the GDPR, the EU AI Act defines “biometric data” as “personal data resulting from specific technical processing relating to the physical, physiological or behavioural characteristics of a natural person”. However, while biometric data do not refer to the psychological characteristics of a person, “biometric identification” includes the automated recognition of “psychological” human features. It is unclear what this could mean in practice. Rec.15 which includes examples of biometric features does not bring clarification in this regard, and whereas methods of physical, physiological, and behavioural biometric identification are reasonably clear (fingerprints, iris scans, facial recognition, mapping body mechanics, etc.), methods of psychological biometric identification are not. As a result, it is difficult to anticipate what categories of information will be treated as psychological biometric identifiers.

□ **“...for the purpose of establishing the identity of a natural person...”** – The concept of identifying a person from biometric data has long been addressed in EU laws – in particular, the GDPR, which treats biometric data as a special category of personal data (Art.9(1) GDPR). As a result, businesses using biometric identification systems qualifying as AI systems will likely need to comply with both the EU AI Act and the GDPR in parallel.

Context and illustrations



Commentary: The concept of “risk”

The concept of risk arises throughout the EU AI Act. AI systems are categorised as having “unacceptable risk”, “high-risk”, “limited risk”, or “minimal risk”. GPAI models are categorised in terms of whether they have “systemic risk”. In each case, the applicable compliance obligations are determined by the risk category into which the relevant system or model falls. In addition, a “risk-based approach” is encouraged (Recs.26 – 27) and providers are required to implement “risk-management systems” (Recs.64 – 65). It is therefore clear that understanding the core concepts of the EU AI Act requires a clear understanding of the concept of “risk”.

The EU AI Act defines “risk” (Art.3(2)) as “the combination of the probability of an occurrence of harm and the severity of that harm”. But in practice, this is very hard to understand. Does it mean that the formula is simply (risk of occurrence) x (severity of harm)? If so, is a highly likely event that will have minimal harm the same as a highly unlikely but potentially catastrophic event? If not, how are we to apply the definition of “risk” in practice? The EU AI Act does not explain. Until clear guidance or case law is provided on this issue, businesses will likely need to make their own risk assessments based on their specificities and be prepared to justify those risk assessments should they be challenged.



Purposive interpretation of the EU AI Act

The challenge of interpreting unclear provisions of EU law is not new. In several GDPR cases in recent years, the CJEU has adopted a “purposive interpretation”. In cases such as *Quadrature du Net (C-511/18)*, *RW v. Post AG (Case C-154/21)* and *FF v. Post AG (Case C-487/21)*, the CJEU resolved complex issues by taking into consideration the purpose of the GDPR (i.e., to ensure a high level of protection of data protection rights of individuals) and opting for the interpretation that best fits that purpose.

Applying the same approach to the EU AI Act would not be quite as straightforward. Whereas the GDPR is essentially focused on a single primary objective (noted above), the EU AI Act has at least two main purposes. Recs.1 – 2 and Art.1 each identify the need to both: (i) ensure a high level of protection for the rights of individuals; **and** (ii) support innovation. These objectives will often be in tension with one another – the greater the protection for individuals, the harder it would be for businesses to lawfully implement innovative AI technologies.

Nevertheless, it is clear that the EU AI Act places the bulk of its emphasis on protecting the rights of individuals, while only Chapter VI deals with innovation in any detail. And even within Chapter VI, the measures in support of innovation repeatedly emphasise the need to protect the rights of individuals. Accordingly, when faced with a question of how best to interpret the provisions of the EU AI Act, courts and regulators will likely adopt the interpretation that is most protective of the rights of individuals.



Commentary: The “design” element of an AI system

As noted above, the definition of an “AI system” requires that a system must (among other things) be “... **designed** to operate with varying levels of autonomy” (emphasis added). But it is uncertain whether the “design” element alone is sufficient – i.e., if a system is “designed” to have autonomy but does not actually have any autonomy in practice, does that system nevertheless meet the definition of an AI system? It is unclear.

The question also cuts the other way – what happens if a system is not “designed” to operate with autonomy but does in practice have autonomy? Is that system potentially an AI system because of the existence of autonomy, even if that autonomy was not intended by the system’s designers? A literal reading of the EU AI Act suggests that such a system would not be an AI system, but courts and regulators may have a different view depending on the peculiarities of each case.

The AI Systems Guidelines refer to two distinct phases in the creation of AI systems: the design/building phase, and the deployment/use phase. The AI Systems Guidelines clarify that the constituent elements of the definition of the term “AI system” are “not required to be present continuously throughout both phases... specific elements may appear at one phase, but may not persist across both phases”. It therefore appears that, even if autonomy was not explicitly part of the design of a system, that system can still be an “AI system” if autonomy emerges in practice during deployment/use. However, this remains an issue that will likely require future clarification from courts and regulators.

Chapter 04

AI literacy

Executive summary

The EU AI Act requires providers (who develop AI systems) and deployers (who use AI systems) to take measures to ensure, to their best extent, that their staff and those dealing with the operation and use of AI systems on their behalf have a sufficient level of AI literacy.

The overarching goal of AI literacy is to ensure that personnel are capable of operating AI systems responsibly and are informed of potential risks and ethical considerations tied to AI's deployment.

It remains to be seen how AI literacy will take shape in practical, real-world applications. The EU AI Act indicates that further guidance will likely be provided in due course. In the interim, businesses should begin developing internal AI literacy, both to ensure employees are equipped to utilise AI systems and to ensure compliance with the EU AI Act.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.4** of the EU AI Act – specifically, outlining what AI literacy might look like in practice and what businesses need to do to ensure AI literacy. This Chapter also includes insights on the relevant definitions in **Art.3**, to the extent that those definitions relate to AI literacy.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model, or that has an AI system or a GPAI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
- **Deployer** (Art.3(4)) – Any organisation using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.
- **AI literacy** (Art.3(56)) – Skills, knowledge, and understanding that allow providers, deployers, and affected persons – taking into account their respective rights and obligations in the context of the EU AI Act – to make an informed deployment of AI systems, as well as to gain awareness about the opportunities and risks of AI and possible harm it can cause.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Understanding AI literacy – The EU AI Act introduces the concept of AI literacy and provides a definition.

An employee is deemed AI literate if they have the skills, knowledge, and understanding to: (i) make an informed deployment of an AI system; and (ii) gain awareness about the opportunities, risks, and possible harm that the relevant AI system may cause.

Whether an employee can make an informed deployment of an AI system will likely depend on the specific context (e.g., what the AI system does and where it will be used). Nevertheless, the wording suggests that employees who are making decisions about AI systems should understand how the AI system works, how it is intended to be used, and how to interpret the AI system's output (Rec.20).

The second aspect of AI literacy is simpler. On a literal reading, employees and other persons acting on behalf of providers and deployers need only have the skills, knowledge, and understanding to gain awareness about the opportunities, risks, and possible harm that an AI system can cause (i.e., those employees and other persons need to know where and how to get more information, but do not necessarily need to have already learned that information in order to satisfy this aspect of AI literacy).

AI literacy in practice – The EU AI Act does not explain what AI literacy looks like in practice, or how businesses can ensure it is achieved. Guidance on AI literacy is expected to evolve gradually, driven by structured support and collaborative initiatives.

Absent any strict guidelines, businesses can focus on developing their employees' AI literacy by providing AI training sessions, workshops, and seminars, or provisionally allocating resources for such initiatives.

The EU AI Act envisions further guidance on AI literacy through coordinated efforts, including:

- **The AI Board support** (Rec.20; Art.66(f)) – The AI Board should support the Commission in promoting AI literacy tools.
- **Codes of conduct** (Rec.20; Art.95(2)(c)) – The EU AI Act encourages the AI Office and EU Member States to take an active role in establishing voluntary codes of conduct to advance and promote AI literacy. It is hoped that these codes of conduct will set consistent standards for AI literacy across industries by serving as frameworks to help organisations implement best practices for ensuring AI literacy.

Once the relevant bodies begin promoting AI literacy and developing voluntary codes of conduct, it is anticipated that businesses will have a clearer understanding of what AI literacy looks like in practice. Meanwhile, businesses that proactively invest in developing AI literacy, or have provisionally allocated resources for developing AI literacy, will be best placed to achieve compliance.

Who needs to be AI literate? The EU AI Act requires providers and deployers to ensure, to their best extent, a sufficient level of AI literacy in two groups of persons:

- **Employees** – Providers and deployers must ensure a sufficient level of AI literacy in their staff.
- **Other persons** – Providers and deployers must ensure a sufficient level of AI literacy in other persons dealing with the operation and use of AI systems on their behalf. For example, if Provider A provided an AI chatbot, other persons might include those that provide technical support to users of the AI chatbot on behalf of Provider A.

Context and illustrations



Commentary: The definition of “AI literacy”

The syntax in the definition of “AI literacy” (Art.3(56)) is unclear. A literal reading of the definition would indicate that the staff of providers and deployers (and other persons dealing with the operation and use of AI systems on their behalf) need only have the skills, knowledge, and understanding to: (i) make an informed deployment of AI systems; and (ii) *gain awareness* about the opportunities and risks of AI, and the possible harm it can cause.

It is advisable that providers and deployers go further than this literal reading and take steps to provide employees and other relevant persons with a sufficient level of AI literacy and suitable opportunities to learn more. Providing education on the *actual* opportunities and risks of the AI system being deployed, and the possible harm that could arise, is unlikely to be a significantly greater burden than ensuring that relevant persons have the skills, knowledge, and understanding to gain such awareness.



Example: Achieving AI literacy

Company X has developed an AI chatbot and implemented it into its workflow. Company Y is engaged by Company X to perform routine maintenance on the AI chatbot.

Company X would need to ensure AI literacy among its staff. Company X would also need to ensure AI literacy at Company Y (as Company Y is a legal person dealing with the operation and use of AI systems on Company X’s behalf).

Until further guidance is provided, Company X should start by assessing the AI knowledge and skills needed across different roles interacting with the AI chatbot. Company X may provide appropriate training, workshops, and seminars explaining measures to be applied to the AI system’s use and/or suitable ways to interpret the AI system’s output (depending on the level of knowledge and skills needed). This could be done through a mixture of hands-on learning and case studies. Company X may also enquire about what Company Y is doing to ensure its staff are AI literate and impose appropriate contractual obligations on Company Y to ensure that the necessary levels of AI literacy are kept up-to-date among Company Y’s staff for the duration of the relationship.

At the very least, Company X is likely to provisionally allocate resources for the purpose of ensuring AI literacy, so that it can effectively deploy such resources once further guidance is provided.

Chapter 05

Prohibited AI practices

Executive summary

The EU AI Act prohibits certain AI practices entirely. Prohibitions came into effect on 2 February 2025, and businesses risk incurring significant penalties for non-compliance (see Chapters 22 and 23 of this Handbook).

The rationale underpinning the specific prohibitions is that the use of AI in particular contexts, and for particular purposes, could result in individuals suffering significant harm. The EU AI Act attempts to mitigate the risk of such harm materialising through the prohibition of specific AI practices.

The list of prohibited practices may change over time following the Commission's annual reviews, which will consider the latest developments in technology.

It is essential that businesses: (i) ensure that they are not engaged in any of the prohibited AI practices; and (ii) monitor the list of prohibited AI practices for changes in the future.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.5** of the EU AI Act – specifically, the list of prohibited AI practices. This Chapter also includes insights on the relevant definitions in **Art.3** as these relate to prohibited AI practices.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Biometric data** (Art.3(34)) – Personal data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of a natural person, such as facial images or dactyloscopic data.
- **Biometric categorisation system** (Art.3(40)) – An AI system for the purpose of assigning natural persons to specific categories on the basis of their biometric data, unless it is ancillary to another commercial service and strictly necessary for objective technical reasons.
- **Remote biometric identification system** (Art.3(41)) – An AI system for the purpose of identifying natural persons, without their active involvement, typically at a distance through the comparison of a person's biometric data with the biometric data contained in a reference database.
- **Real-time remote biometric identification systems** (Art.3(42)) – A remote biometric identification system, whereby the capturing of biometric data, the comparison, and the identification all occur without a significant delay, comprising not only instant identification, but also limited short delays in order to avoid circumvention.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Prohibited AI practices: Businesses – Art.5 of the EU AI Act prohibits five specific AI practices that are particularly relevant for businesses:

- **Subliminal, manipulative or deceptive techniques** (Rec.29; Art.5(1)(a)) – The EU AI Act prohibits AI systems that deploy subliminal, purposefully manipulative, or deceptive techniques with the objective or effect of materially distorting behaviour in a manner which causes (or is reasonably likely to cause) significant harm.

It is worth noting that the AI system does not need to set out to materially distort behaviour to be prohibited; it is sufficient that the AI system has such an effect for this prohibition to apply. This prohibition also contains notable qualifiers (i.e., distortion of behaviour must

be **material**, and must cause (or be reasonably likely to cause) **significant** harm). These qualifiers will help ensure this prohibition does not apply to common and legitimate commercial practices in the field of advertising that are otherwise compliant with applicable law (Rec.29).

- **Exploiting vulnerabilities** (Rec.29; Art.5(1)(b)) – The EU AI Act prohibits AI systems that exploit any vulnerabilities of a natural person(s) due to age, disability, or social/economic situation with the objective or effect of materially distorting behaviour in a manner which causes (or is reasonably likely to cause) significant harm.

Again, it is not necessary that the AI system intends to materially distort behaviour to be prohibited; the prohibition will apply if the AI system has this effect. This prohibition contains the same qualifiers as the prohibition concerning subliminal, manipulative, or deceptive techniques (see above). Note: This prohibition (and the prohibition described above) should not apply to lawful medical practices carried out in accordance with applicable medical standards (Rec.29).

- **Facial recognition via untargeted scraping** (Rec.43; Art.5(1)(e)) – The EU AI Act prohibits AI systems that create or expand facial recognition databases via the untargeted scraping of facial images from the internet or CCTV footage.

On a literal reading, this prohibition applies only to “*untargeted*” scraping of facial images from the internet or CCTV footage, not to targeted scraping.

- **Inference of emotions in the workplace and educational institutions** (Rec.44; Art.5(1)(f)) – The EU AI Act prohibits AI systems that infer emotions of a natural person in the workplace and educational institutions, except for medical or safety reasons (such as systems intended for therapeutical use (Rec.44)).

The EU AI Act distinguishes between internal emotions (e.g., happiness, sadness, anger, surprise, or disgust) and physical states or expressions (e.g., pain, fatigue, readily apparent expressions, gestures, or movements). Inference of physical states or expressions are not caught by this prohibition (Rec.18).

- **Biometric categorisation** (Rec.30; Art.5(1)(g)) – The EU AI Act prohibits AI systems that use biometric data to deduce or infer race, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation.

AI systems that use biometric data to infer ethnicity, health data, or genetic data will be high-risk AI systems and are not prohibited by Art.5 (Rec.54). Please refer to Chapter 6 of this Handbook for more information on the classification of AI systems as high-risk.

Prohibited AI practices: Governments – Art.5 of the EU AI Act prohibits two specific AI practices that are particularly relevant for governments:

- **Social scoring in certain use cases** (Rec.31; Art.5(1)(c)) – The EU AI Act prohibits AI systems that apply social scoring to individuals/groups of individuals (i.e., evaluating or classifying individuals over time according to, amongst other things, social behaviour and personal/personality characteristics) if the social scores lead to detrimental or unfavourable treatment of individuals in either or both of the following ways: (i) in social contexts unrelated to the contexts in which the data was originally generated or collected; and/or (ii) in ways that are unjustified or disproportionate to their social behaviour or its gravity.
- **Predictive policing** (Rec.42; Art.5(1)(d)) – The EU AI Act prohibits AI systems from making risk assessments to assess or predict the likelihood of an individual committing a criminal offence, based solely on profiling or assessing the relevant individual’s personality traits and characteristics. This prohibition does not apply when AI systems are used to support human assessment of an individual’s criminal activity based on objective and verifiable facts directly linked to a criminal activity.

Prohibited AI practices: Law enforcement – The EU AI Act also prohibits AI systems that use “real-time” remote biometric identification systems (RBIS) in publicly accessible spaces for law enforcement purposes, unless an exception applies (Rec.32; Art.5(1)(h)). Even if an exception applies, the use remains subject to certain conditions and reporting obligations. For further details of the applicable conditions and reporting obligations, please see Recs.33 – 35 and Arts.5(2) – (7).

Exceptions – The use of “real-time” RBIS in publicly accessible spaces for law enforcement purposes is permitted if it is strictly necessary for:

- The targeted search for specific victims of abduction, trafficking, or sexual exploitation, or the search for missing persons (Rec.33; Art.5(1)(h)(i)).
- The prevention of a specific, substantial, and imminent threat to life or physical safety, or a genuine and present/foreseeable threat of a terrorist attack (Rec.33; Art.5(1)(h)(ii)).
- The localisation, identification, or prosecution of an individual suspected of having committed a criminal.

Future review of prohibited AI practices – The EU AI Act allows for amendments to be made to the list of prohibited AI practices over time, taking into account relevant evidence and developments in technology (Art.112). Businesses will need to remain up-to-date with the current list of prohibited AI practices.

Once a year from 1 August 2024, the Commission shall assess the need to amend the list of prohibited AI practices contained in Art.5. The findings of that assessment must be submitted to the European Parliament and the Council (Art.112(1)). If necessary, the Commission will also submit appropriate proposals to amend the EU AI Act (Art.112(10)).

Accordingly, the list of prohibited AI practices may change over time, and businesses will need to monitor any such changes in order to ensure continued compliance with the EU AI Act.

Context and illustrations



Commentary: “Objective” or “effect”

The scope of the prohibitions contained in Art.5(1)(a) and (b) is broad. Both prohibitions are expressed as applying to AI systems that have “*the objective, or the effect of materially distorting*” the behaviour of certain persons, either by deploying subliminal, manipulative, or deceptive techniques, or exploiting certain vulnerabilities. Rec.29 similarly states that “*it is not necessary for the provider or the deployer to have the intention to cause significant harm, provided that such harm results from the manipulative or exploitative AI-enabled practices*”.

As noted above, this means that an AI system does not need to set out to materially distort behaviour to be prohibited; it is sufficient that the AI system has such an effect for the application of these prohibitions.

This wording suggests that businesses will need to invest in resources to ensure compliance with Art.5(1)(a) and (b), both before deployment and thereafter. For example, businesses should: (i) engage with legal counsel at an early stage to clarify and document an AI system’s intended and anticipated effects; (ii) ensure alignment between technical and legal teams on an AI system’s intended and anticipated effects; and (iii) be prepared to react on an ongoing basis if the AI system produces unanticipated effects that fall within the scope of Art.5(1)(a) or (b) once deployed.



Example: Prohibited vs. non-prohibited AI practices

Company X is a shoe company, which has created two billboard ads and is unsure which is more effective. Company Y develops an AI system that can scan facial images and infer emotional reactions. Company X places Company Y’s AI system in its billboards to gauge the emotional reactions of those encountering the ads for the purposes of determining which of the two is more popular.

This does not appear to be a prohibited AI practice for the following reasons:

- The AI system does not deploy subliminal, manipulative, or deceptive techniques that materially distort behaviour in a manner which causes (or is reasonably likely to cause) significant harm;
- The AI system does not exploit vulnerabilities of natural person(s);
- The scanning of facial images is arguably targeted (i.e., it only scans the face of those who pass the billboard, so it has targeted geographic scope);
- The AI system does not infer emotional reactions in the workplace or educational institutions; and
- The biometric data is not used to infer race, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation.

Whereas, if Company X used Company Y’s AI system to gauge the reaction of passers-by to two different **political** ads, this would likely contravene Art.5(1)(g) and amount to a prohibited AI practice. Such use of biometric data would likely be inferring the political opinions of those who pass by the billboards.



Commentary: “Common and legitimate commercial practices”

There is some ambiguity as to what “*common and legitimate commercial practices*” might benefit from the carve-outs in Art.5(1)(a) and (b) in light of Rec.29.

Art.5(1)(a) and (b) prohibit AI systems that deploy subliminal, purposefully manipulative, or deceptive techniques, or exploit certain vulnerabilities, with the objective or effect of materially distorting behaviour in a manner which causes (or is reasonably likely to cause) **significant** harm.

As noted above, Rec.29 suggests that these prohibitions are not concerned with “*common and legitimate commercial practices*” that comply with applicable law. Such AI systems might include, for example, those used for the purpose of personalised advertising or dynamic pricing, as these may reasonably be considered “*common and legitimate commercial practices*”.

However, neither of the prohibitions contained in Arts.5(1)(a) nor 5(1)(b) reflects the wording of Rec.29; instead, the mention of “*common and legitimate commercial practices*” is contained in a non-operative provision of the EU AI Act. Consequently, any defence of an AI system on the basis that it is a “*common and legitimate commercial practice*” may be difficult to maintain, particularly if it materially distorts behaviour in a manner which causes (or is reasonably likely to cause) significant harm.

Chapter 06

Classification of AI systems as “high-risk”

Executive summary

The EU AI Act adopts a risk-based approach, in which AI systems that are categorised as “high-risk” are subject to stringent requirements. AI systems are deemed to be “high-risk” on the basis of the categories into which they fall, rather than a fact-based analysis of the actual level of real-world risk associated with each AI system.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.6** and **Annexes I and III** of the EU AI Act – specifically, the grounds on which an AI system will be classified as a high-risk AI system for the purposes of the EU AI Act.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **AI system** (Rec.12; Art.3(1)) – A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment. For explicit or implicit objectives, an AI system infers from the input it receives how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.
- **Risk** (Art.3(2)) – The combination of the probability of an occurrence of harm and the severity of that harm.
- **Safety component** (Art.3(14)) – A component of a product or of an AI system which fulfils a safety function for that product or AI system. The failure or malfunctioning of a safety component is something which endangers the health and safety of persons or property.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Risk-based approach – The EU AI Act imposes different sets of rules on AI systems. This ranges from minimal risk at the lowest end, increasing to limited risk, then to high-risk, and finally to the maximum level of unacceptable risk. AI systems that carry unacceptable risks are prohibited (see Chapter 5). Similarly, high-risk AI systems are subject to stringent requirements, as discussed in Chapters 7 and 8.

High-risk AI systems are divided into two sub-categories that are subject to different rules and requirements. As a result, businesses need to ensure that they review each of their AI systems and understand which category each AI system falls into.

Businesses need to understand which rules apply to their AI systems. To figure this out, they should apply the following tests:

The “not prohibited” test: There is some overlap between the prohibitions in Art.5 (discussed in Chapter 5) and the definition of high-risk AI systems. As a result, the first question each business should ask is whether the AI system in question is prohibited. If the AI system is prohibited, there is no need to consider whether it is high-risk.

The “safety-critical AI systems” test: Some AI systems may carry the risk of causing an adverse impact on health and safety, when such systems are part of, or are used as, safety-critical products. The safety-critical AI systems test, as established in Art.6(1), helps to identify these risky systems in two steps (Recs.47, 50, and 51):

The first step is to determine whether the AI system falls under one of the lists of EU safety laws listed in the two sections of **Annex I** to the AI Act (Art.6(1)(a)).

Legislation under Section A relates to the following:

- **Appliances burning gaseous fuels** (for example, AI-controlled efficiency optimisation systems for hot water boilers may fall within this sub-category).
- **Cableway installations** (for example, AI-driven safety monitoring systems for cableway operations may fall within this sub-category).

- **Equipment and protective systems intended for use in potentially explosive atmospheres** (for example, AI-driven mining robots may fall within this sub-category).
- **In vitro diagnostic medical devices** (for example, AI-driven diagnostic tools and AI-powered robotic surgical systems used in in vitro devices may fall under this sub-category).
- **Lifts and safety components for lifts** (for example, AI-driven safety monitoring systems and predictive maintenance systems in lifts may fall within this sub-category).
- **Machinery** (for example, AI systems integrated into industrial robots or other automated machinery may fall within this sub-category).
- **Medical devices** (for example, AI-driven diagnostic tools and AI-powered robotic surgical systems may fall within this sub-category).
- **Personal protective equipment** (for example, AI-powered wearable safety gear may fall within this sub-category).
- **Pressure equipment** (for example, AI systems that monitor and control pressure equipment may fall within this sub-category).
- **Radio equipment** (for example, electronic devices with AI capabilities that transmit and/or receive radio signals (e.g., Wi-Fi and Bluetooth) may fall within this sub-category).
- **Recreational craft and personal watercraft** (for example, smart sensors in boats may fall within this sub-category).
- **Toys** (for example, smart toys with AI capabilities and AI-driven gaming systems may fall within this sub-category).

Legislation under Section B relates to the following:

- **Agricultural and forestry vehicles** (for example, AI-driven harvesting robots may fall within this sub-category).
- **Civil aviation security** (for example, AI-powered civil aviation threat detection systems may fall within this sub-category).
- **Marine equipment** (for example, smart sensors for marine safety may fall within this sub-category).
- **Motor vehicles and their trailers, as well as systems, components, and separate technical units intended for such vehicles** (for example, automated braking systems could fall under this sub-category).
- **Rail systems** (for example, predictive maintenance of rail infrastructure may fall within this sub-category).
- **Two- or three-wheel vehicles and quadricycles** (for example, AI systems for autonomous driving may fall within this sub-category).
- **Unmanned aircraft and their engines, propellers, parts, and equipment to control them remotely** (for example, AI-driven controls for drones and air traffic may fall within this sub-category).

AI systems that fall into Section B above are exempted from the majority of the requirements of the EU AI Act (Art. 2(2)).

The second step is to determine whether the AI system requires a **third-party conformity assessment** under the laws outlined above, in order to be placed on the market or put into service (Art. 6(1)(b)). If the AI system does not require a third-party conformity assessment under the laws outlined above, then that AI system is not deemed to be “high-risk” under this “safety-critical AI systems” test. However, if the AI system requires a third-party conformity assessment under these laws, then it is deemed to be “high-risk” for the purposes of the EU AI Act and is subject to the requirements set out in Chapters 7 and 8.

The “high-risk categories” test: This test also involves two steps (Art.6(2) and (3); Annex III):

The first step is to determine whether the AI system is intended to be used for any of the purposes listed in Annex III:

- **Biometrics** – Remote biometric identification systems, AI systems that are intended to be used for biometric categorisation, and AI systems intended to be used for emotion recognition fall within this sub-category.
- **Critical infrastructure** – AI systems that are intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or in the supply of water, gas, heating, or electricity fall within this sub-category.
- **Educational and vocational training** – AI systems that are intended to be used to determine access or admissions to educational and vocational training, to evaluate learning outcomes, to assess the appropriate level of education that an individual may receive or access, or to monitor and detect prohibited student behaviours fall within this sub-category.
- **Employment, workers’ management, and access to self-employment** – AI systems that are intended to be used to recruit individuals, or to make decisions affecting or relating to employment fall within this sub-category.
- **Access to and enjoyment of essential private services and essential public services and benefits** – AI systems that are intended to be used to assess creditworthiness, or for life and health insurance-related assessments or pricing, and for responding to requests for emergency response or healthcare services fall within this sub-category. AI systems that are intended to be used by public authorities to assess the eligibility of individuals for essential public assistance (e.g., benefits and/or services) will also fall within this sub-category.
- **Law enforcement** – AI systems that are intended to be used by/on behalf of/in support of law enforcement authorities to assess the risk of an individual becoming a victim/offender/re-offender of a criminal offence, or to evaluate the reliability of evidence, or to assess personality traits and characteristics or past criminal behaviour of individuals, or to detect, investigate, or prosecute criminal offences will fall within this sub-category. AI systems intended to be used as polygraph or similar will also fall within this sub-category.
- **Migration, asylum, and border control management** – AI systems intended to be used by/on behalf of/in support of competent public authorities for certain purposes related to migration, asylum, and border control management will fall within this sub-category.
- **Administration of justice** – AI systems that are intended to be used by/on behalf of judicial authorities for researching and interpreting facts and the law, or for applying the law to a concrete set of facts, or that are intended to be used in a comparable way in alternative dispute resolution, will fall within this sub-category.
- **Administration of democratic processes** – AI systems that are intended to influence the outcome of an election or referendum, or the voting behaviour of individuals in exercising their right to vote in an election or referendum, will fall within this sub-category. Note that this sub-category does not include AI systems used merely to organise, optimise, or structure political campaigns from an administrative or logistical point of view.

The second step to the “high-risk categories” test is to consider whether the AI system might be exempt from “high-risk” status on the basis that it does **not pose a significant risk** of harm to the health, safety, or fundamental rights of individuals (Rec.53 and Art. 6(3)). AI systems that are intended to be used for the following purposes are not deemed high-risk under Art.6(2):

- Any AI system that is intended to **perform a narrow procedural task** (such as transforming unstructured data into structured data, classifying documents into categories, or detecting duplicates among a large number of applications (Rec.53)).
- Any AI system that is intended to **improve the result of a previously completed human activity** (e.g., improving the language used in previously drafted documents, in relation to professional tone, academic style of language, or aligning text with a particular brand messaging).
- Any AI system that is intended to **detect decision-making patterns or deviations** from prior decision-making patterns and is not meant to replace or influence a previously completed human assessment (e.g., an AI system that can be used to check whether decision makers have deviated from the standard decision-making patterns in order to identify potential inconsistencies or anomalies).
- Any AI system intended to perform a mere **preparatory task for a risk assessment** (e.g., enabling smart solutions for file handling that includes functions such as indexing, searching, text and speech processing, or linking data to other data sources, or AI systems for machine translations).

However, notwithstanding the preceding bullets, any AI system that is used for a purpose listed in Annex III and is also **used to profile individuals** will always be considered high-risk.

To rely on an exemption under Art. 6(3), businesses will be expected to retain details of the assessment they carried out in order to determine that an exemption was available, prior to the placing of the AI system on the market, and/or putting it into service. Businesses will also be required to register any exempt AI system in the EU database for high-risk AI systems (Art.6(4)).

Context and illustrations



Practical tip: Identifying high-risk AI systems

The EU AI Act requires the Commission to publish guidelines specifying the practical implementation of Art.6. This includes a comprehensive list of practical examples of high-risk and non-high-risk AI systems (Art.6(5)). The Commission has until 2 February 2026 to publish these guidelines. Until then, businesses will have to consider their exposure to Chapter 3 of the EU AI Act with limited regulatory guidance.

As a first step, businesses should build a clear understanding of how they are using AI systems, and for what purposes. A clear understanding of these two factors can serve as the basis for a comprehensive assessment of the businesses' use of AI more generally, i.e., to establish which AI systems fall within the scope of Art.6 and to what extent any exemptions are available under the EU AI Act.



Commentary: Purposive interpretation of the concept of high-risk AI systems

In its pursuit of a technology-neutral and future-proof AI legislation, the EU AI Act's approach to explaining core concepts – including that of high-risk AI systems – is marred by complexity and uncertainty. The ambiguities apparent in Art.6 and Annexes I and III leave ample room for grey areas when it comes to the classification of high-risk AI systems. In particular, an AI system that poses a high level of real-world risk but does not fall into any of the categories in Arts.6(1) or 6(2) will not be a “*high-risk AI system*”; whereas an AI system that does fall into those categories will be a “*high-risk AI system*” even if the real-world level of risk is comparatively low.

In the enforcement context, courts and regulators in the EU may resort to a highly context-specific interpretation, driven by what each respective court views as the purpose of the regulation, as well as their perception of risk.

However, there are several Recitals in the EU AI Act relating to high-risk AI systems, in particular Recs.47 – 63, which at least provide businesses with indications of the types of AI systems that might fall into one of the categories.



Commentary: Intended purpose and unintended use

When assessing exposure to Art.6, businesses should bear in mind that the classification of an AI system as high-risk generally depends on the “intended purpose” of the AI system in question.

The EU AI Act defines the term “*intended purpose*” as the use for which an AI system is **intended** by the provider, including (but not limited to) the specific context and conditions of use. This is something that is specified in the information that is supplied by the provider in the instructions for use, or in promotional or sales material and statements, as well as in the technical documentation (Art.3(12)).

Businesses will need to account for situations in which an AI system is used for a purpose that was not intended by the provider, and may be required to address risks that could foreseeably arise from such uses. Businesses will also need to account for situations in which an AI system has more than one intended purpose.

Chapter 07

Requirements for high-risk AI systems

Executive summary

The EU AI Act establishes mandatory requirements for high-risk AI systems. Key requirements include continuous risk management processes, ensuring the quality and bias mitigation of data sets, and preparing detailed technical documentation. High-risk AI systems are also required to have automatic event logging for traceability, be accompanied by clear and comprehensive user information, and incorporate effective human oversight measures to minimise risks. Additionally, these systems should be designed to maintain high levels of accuracy, robustness, and cybersecurity throughout their lifecycle.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.8 to 15** of the EU AI Act which set out mandatory requirements regarding risk management systems (**Art.9**), data and data governance (**Art.10**), technical documentation (**Art.11**), recordkeeping (Art.12), transparency and user information (**Art.13**), human oversight (**Art.14**) and accuracy, robustness, and cybersecurity (**Art.15**).

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Risk** (Art.3(2)) – The combination of the probability of an occurrence of harm and the severity of that harm.
- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model, or that has an AI system or a GPAI model developed and places it on the market, or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
- **Deployer** (Art.3(4)) – Any organisation using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.
- **Intended purpose** (Art.3(12)) – The use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Compliance with the requirements (Art.8) – Providers of high-risk AI-systems should comply with the requirements laid down in Arts.9 to 15.

- **Compliance requirements in general** – Providers of high-risk AI systems are required to satisfy the compliance obligations set out in this Chapter. They must take into account the intended purpose of the relevant high-risk AI system and the state of the art of AI technologies. In addition, providers of high-risk AI systems need to ensure compliance with the EU harmonisation legislation listed in Section A of Annex I of the EU AI Act.

Risk management systems (Art.9) – Providers must establish, implement, document, and maintain a risk management system for each high-risk AI system they provide.

- **Nature of risk management systems** – The risk management system needs to adapt over time. It should be thought of as an ongoing, iterative process throughout the AI system’s lifecycle. This process should include: (a) identifying and analysing risks to health, safety, and fundamental rights; (b) evaluating risks that arise during the use of the AI system; (c) evaluating other potential risks based on post-market monitoring system data; and (d) adopting appropriate risk management measures (Art.9(2)-(3)).
- **Purpose of risk management systems** – The goals of a risk management system are to: (a) reduce or eliminate identified risks (as far as technically feasible); (b) implement appropriate mitigation where risks cannot be eliminated; (c) provide transparency information in accordance with Art.13; and (d) where appropriate, provide training to deployers of the relevant AI systems (Art.9(5)).
- **Testing obligations** – Providers must test their high-risk AI systems in order to identify the most appropriate risk management measures, and to ensure that those AI systems are performing as intended (Art.9(6)-(8)).

Data management and governance (Art.10) – High-risk AI systems should be developed using high-quality data sets for training, validation, and testing.

- **Data governance** – Providers of high-risk AI systems must ensure that training, validation, and testing data sets are appropriate for the intended purpose of the high-risk AI system (Art.10(1)-(2)). Providers must implement appropriate data governance and management practices to address issues including design choices, data collection processes, data sources, data preparation, assumptions, and the availability, quantity, and suitability of data sets (Art.10(2)(a)-(e)).
- **Detecting, preventing, and mitigating possible biases** – Providers must ensure that data governance and management practices include appropriate measures to detect, prevent, and mitigate possible biases that are likely to affect health and safety, fundamental rights, or lead to unlawful discrimination (Art.10(2)(f)-(g)). Providers are also required to identify and mitigate data gaps or shortcomings in relation to the foregoing issues (Art.10(2)(h)).
- **Processing SCD for bias correction purposes** – Providers are “*exceptionally*” permitted to process special categories of personal data (SCD) for the purposes of bias correction, provided that they implement appropriate safeguards for affected data subjects (Art.10(5)). In addition to satisfying the requirements of the GDPR and other relevant EU laws, providers must ensure that: (i) processing SCD is the only effective way to detect and correct bias; (ii) the SCD are appropriately protected and are not further shared; (iii) the SCD are deleted once the bias has been corrected; and (iv) appropriate records of processing activities are maintained recording the objectives and necessity of the processing.

Technical documentation (Art.11) – Providers must draw up technical documentation of high-risk AI systems before those systems are placed on the market or put into service.

- **Purpose of technical documentation** – Technical documentation must be drawn up in advance. It must demonstrate that the high-risk AI system meets the requirements outlined in this Chapter (Art.11(1)). It must contain the information specified in Annex IV of the EU AI Act. In effect, this requires businesses to create and maintain detailed transparency information about each high-risk AI system they provide.
- **Keeping technical documentation up-to-date** – The technical documentation noted above needs to be kept up-to-date. The Commission may amend Annex IV (i.e., the list of information that the technical documentation needs to address) (Art.11(1)). As a result, providers need to keep track of both the changes they make to their high-risk AI systems and also the changes the Commission introduces to Annex IV.

Recordkeeping (Art.12) – Providers of high-risk AI systems need to implement automated event logging.

- **Traceability** – The EU AI Act requires providers to ensure that high-risk AI systems have automated event logging to enable a level of traceability appropriate to the system's intended purpose (Art.12(1)). This event logging capability should include:
 - Identifying risks to health, safety, or fundamental rights (Arts.12(2)(a) and 79(1)) (see Chapter 19).
 - Facilitation of post-market monitoring (Arts.12(2)(b) and 72) (see Chapter 19).
 - Monitoring of the system's operation (Arts.12(2)(c) and 26(5)) (see Chapter 8).

These logs need to be retained for at least six months (Art.19(1)) (see Chapter 8).

For remote biometric identification systems, the logging system needs to include additional detail regarding recording periods and data sources (Art.12(3)).

Transparency information (Art.13) – High-risk AI systems should be designed so that their workings are transparent and can be understood by deployers.

- **Transparency and provision of information to deployers** – Providers must ensure that high-risk AI systems are accompanied by instructions for deployers. At a minimum, this should include the identity and contact details of the provider; the characteristics, capabilities, and limitations of performance of the high-risk AI system (e.g., including its intended purpose); the changes to the system; human oversight measures referred to in Art.14 (see below); information about required hardware; and expected lifetime and maintenance and care measures, as well as a description of the mechanisms included in the high-risk AI system that allow deployers to properly collect, store, and interpret the logs in accordance with Art.12 (see above). Additional transparency obligations that apply in relation to AI systems are discussed in Chapters 11 (in relation to certain AI systems) and 13 (in relation to GPAI models).

Human oversight (Art.14) – High-risk AI systems should be designed and developed in such a way that they can be effectively overseen by a human.

- **Goal of human oversight** – The goal of human oversight of high-risk AI systems is to prevent or minimise the risks to health, safety, or fundamental rights (Art.14(2)).
- **Achieving human oversight** – Human oversight of high-risk AI systems should be achieved through: (a) built-in measures added to the high-risk AI system prior to launch; and/or (b) measures identified by the provider after launch that can be implemented by deployers (Art.14(3)). Providers need to provide their high-risk AI systems to deployers in such a way that the human overseeing the system can understand its capabilities and limitations, detect and address issues, avoid over-reliance/automation bias, correctly interpret its output, decide not to use it, or stop its operation.

Accuracy, robustness, and cybersecurity (Art.15)

– High-risk AI systems should be designed and built to achieve an appropriate level of accuracy, robustness, and cybersecurity.

- **Explaining accuracy** – Providers need to explain the levels of accuracy that their high-risk AI systems achieve. The Commission is required to work with the industry to determine how best to measure levels of accuracy, robustness, and other relevant performance metrics (Art.15(2)). In the interim, providers should consider how best to explain the levels of accuracy and robustness that their high-risk AI systems should be expected to achieve in normal use – this information needs to be explained in the instructions provided to deployers (Arts.13(3)(b)(ii) and 15(3)).
- **Resilience against errors** – The EU AI Act obliges providers to ensure that their high-risk AI systems are “*as resilient as possible*” against errors, faults, or inconsistencies (especially in relation to interactions with individuals or other systems). Providers are required to take technical and organisational measures to achieve this goal (e.g., testing, backups, disaster recovery plans, etc.) (Art.15(4)). However, the EU AI Act does not explain the requirement to be “*as resilient as possible*”. It seems, on the face of it, that this standard is impossible to satisfy in the literal sense. No matter how good the existing resilience measures may be, there is always more that could be done, even if doing so would only have a very minor effect. Therefore, the requirement to be “*as resilient as possible*” can never be completely satisfied, since it is always “*possible*” to do more.

- **Cybersecurity** – Providers need to take appropriate measures to ensure that their high-risk AI systems are resilient against cyberattacks, including measures to prevent, detect, respond to, resolve, and control such attacks. The level of cybersecurity achieved needs to be appropriate to the relevant circumstances and the risks (Art.15(5)).

Context and illustrations



Commentary: Achieving pragmatic compliance with the requirements for high-risk AI systems

As noted above in this Chapter, Arts.8 to 15 of the EU AI Act set out a series of complex obligations that providers of high-risk AI systems are required to satisfy.

While the means of achieving such compliance are not always clear, it makes sense to implement internal measures to help achieve compliance as far as possible.

To that end, providers of high-risk AI systems should consider setting up internal governance structures that set out clear rules and checklists for employees and management teams to follow, in particular: (i) setting up a risk management system to review all high-risk AI systems on an ongoing basis and identify, evaluate, and mitigate risks; (ii) implementing appropriate internal governance procedures (e.g., data quality guidelines, anti-bias guidelines, etc.); (iii) keeping up-to-date comprehensive technical documentation about each high-risk AI system; (iv) establishing automated logging procedures and systems to maintain the relevant logs; (v) creating an AI human oversight procedure to ensure that humans can understand, interpret, and intervene in the system's operations; (vi) preparing an accuracy and robustness procedure setting out the steps that need to be followed in the design and operation of the system; (vii) ensuring that existing cybersecurity policies are updated to address the requirements of the EU AI Act; and (viii) keeping a close eye on both the evolution of the relevant high-risk AI systems over time, and developments and guidance coming from the Commission and regulators.



Analysis: "Intended purpose"

Arts.8 –15 of the EU AI Act repeatedly refer to the concept of the "*intended purpose*" of a high-risk AI system. That concept is defined in Art.3(12) as "*the use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation*".

It is therefore essential for each provider to: (i) ensure that there is internal clarity on the intended use case(s) of each high-risk AI system (e.g., by maintaining clear internal records of what each system is designed to do, the conditions in which it should be used, and any purposes for which it should not be used); (ii) ensure that all promotional materials are consistent with the intended use case(s); and (iii) ensure that the relevant intended use case(s) for each high-risk AI system are fully and consistently reflected in the technical documentation (Art.11), recordkeeping (Art.12), and transparency materials (Art.13) noted above in this Chapter.

Without these measures, there is a risk that a court or regulator may adopt a different interpretation of a high-risk AI system's "*intended purpose*" to the interpretation intended by the provider, making it significantly harder for the provider to demonstrate compliance.



Commentary: Intended purpose and unintended use

When assessing exposure to Art.6, businesses should bear in mind that the classification of an AI system as high-risk generally depends on the "*intended purpose*" of the AI system in question.

The EU AI Act defines the term "*intended purpose*" as the use for which an AI system is **intended** by the provider, including (but not limited to) the specific context and conditions of use. This is something that is specified in the information that is supplied by the provider in the instructions for use, or in promotional or sales material and statements, as well as in the technical documentation (Art.3(12)).

Businesses will need to account for situations in which an AI system is used for a purpose that was not intended by the provider, and may be required to address risks that could foreseeably arise from such uses. Businesses will also need to account for situations in which an AI system has more than one intended purpose.

Chapter 08

Obligations relating to high-risk AI systems

Executive summary

The EU AI Act imposes stringent obligations on providers and deployers of high-risk AI systems, and somewhat less burdensome obligations on other relevant parties (e.g., authorised representatives, importers, and distributors). Providers, deployers, and other parties involved each have different obligations with respect to high-risk AI systems. In addition, distributors, importers, and deployers are considered providers in certain cases.

Relevant sections of the EU AI Act

This Chapter focuses on **Section 3** of the EU AI Act, which sets out compliance obligations of providers, deployers, and other parties in relation to high-risk AI systems. The specific obligations are laid out in **Arts.16 to 27**.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model or that has an AI system or a GPAI model developed, and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
 - **Deployer** (Art.3(4)) – Any organisation using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.
 - **Authorised representative** (Art.3(5)) – Any organisation located or established in the EEA who has received and accepted a written mandate from a provider of an AI system or a GPAI model to, respectively, perform and carry out on its behalf the obligations and procedures established by the EU AI Act.
 - **Importer** (Art.3(6)) – Any organisation located or established in the EEA that places on the market an AI system that bears the name or trademark of any organisation outside the EEA.
 - **Distributor** (Art.3(7)) – Any organisation in the supply chain, other than the provider or the importer, that makes an AI system available on the EEA market.
- A full list of defined terms can be found in the [Glossary](#).

Analysis

General obligations of providers (Arts.16 to 21) –

Arts.16 to 21 set out specific obligations that providers of high-risk AI systems must fulfil. Art.16 contains providers' primary obligations, while Arts.17 to 21 set out specifications of some of the obligations referred to in Art.16.

- **Obligations of providers (Art.16)** – Art.16 requires that providers of high-risk AI systems must: ensure that their high-risk AI systems comply with the requirements set out in Arts.8 to 15 (see Chapter 7) (Art.16(a)); indicate their name and contact details on the AI system or its packaging (Art.16(b)); have a quality management system in place (Arts.16(c) and 17); keep the necessary documentation (Arts.16(d) and 18); keep automatically generated logs (Arts.16(e) and 19); undergo the relevant conformity assessment procedure (Arts.16(f) and 43); draw up an EU declaration of conformity (Arts.16(g) and 47); affix the CE marking to the AI system (Arts.16(h) and 48); comply with the registration obligations (Arts.16(i) and 49); take necessary corrective actions (Arts.16(j) and 20); cooperate with requests from national competent authorities to demonstrate compliance with Arts.8 to 15 (see Chapter 7) (Arts.16(k) and 21); and comply with accessibility requirements under EU law (Art.16(l)).

- **Quality management system (Art.17)** – Providers of high-risk AI systems must put a quality management system in place that ensures compliance with the EU AI Act. The AI system should be documented, and must include at least the following: a strategy for regulatory compliance, such as conformity assessment procedures; techniques for the design, development, and quality control of the AI system; examination and testing procedures; technical specifications; systems for data management; a risk management system; a post-market monitoring system pursuant to Art.72; procedures for reporting serious incidents pursuant to Art.73; process for handling communications with relevant authorities; a system for recordkeeping; and resource management, as well as an accountability framework setting out the management and staff responsibilities regarding all aforementioned aspects of the quality management system. Art.17(2) states that implementation of the quality management system should be “*proportionate to the size of the provider’s organisation*”. Rec.146 indicates that this is intended to allow “*microenterprises*” to implement a quality management system in a “*simplified manner*”, reducing administrative burdens and costs. It is not yet entirely clear how this will work in practice. The Recital indicates that the Commission will develop guidelines on this issue. Providers that are subject to EU rules on financial institutions are subject to additional obligations – see below.
- **Documentation keeping (Art.18)** – Providers must maintain, and on request provide, national competent authorities with the technical documentation they are required to maintain under Art.11; documentation concerning the quality management system (Art.17); documentation concerning the changes approved by notified bodies (see Chapter 9); the decisions and other documents issued by the notified bodies (see Chapter 9); and the EU declaration of conformity that the provider is required to maintain pursuant to Art.47. These items must be maintained for a period ending ten years after the AI system has been placed on the market or put into service. Providers that are subject to EU rules on financial institutions are subject to additional obligations – see below.
- **Automatically generated logs (Art.19)** – Providers of high-risk AI systems must keep the logs that are automatically generated over the lifetime of their systems (Art.12(1)), to the extent such logs are under their control. These should be kept for a period of at least six months, except where a longer retention period is required by EU law, or applicable Member State law. Providers that are subject to EU rules on financial institutions are subject to additional obligations – see below.
- **Corrective actions and duty of information (Art.20)** – If providers consider (or have reason to consider) that a high-risk AI system that they have placed on the market or put into service does not conform with the EU AI Act, they must “*immediately*” take the necessary corrective actions to bring that system into conformity, to withdraw it, to disable it, or to recall it. They must also inform the distributors of the AI system and, where applicable, the deployers, the authorised representative, and importers accordingly. If the high-risk AI system presents a risk to health, safety, or fundamental rights of persons, and the provider becomes aware of that risk, it must immediately investigate the causes, inform the competent market surveillance authority and, where applicable, the notified body that issued the conformity certificate for that high-risk AI system. It is not entirely clear at what point a provider is deemed to have reason to consider that a high-risk AI system is not in conformity, but Rec.115 indicates that this is an issue that providers would be expected to detect through the post-market monitoring systems they are required to establish in accordance with Art.72 (see Chapter 19).
- **Cooperation with competent authorities (Art.21)** – Providers of high-risk AI systems must, upon a reasoned request by a competent authority, provide all the information and documentation necessary to demonstrate the conformity of the high-risk AI system with the requirements set out in Arts.8 to 15 (see Chapter 7).

Obligations of providers established outside the EEA (Art.22) – Providers established outside the EEA must, by written mandate, appoint an authorised representative in the EEA.

The authorised representative will be required to carry out the following tasks (which will be specified in a mandate received from the provider): Verify that the EU declaration of conformity and necessary technical documentation have been drawn up; verify that an appropriate conformity assessment procedure has been carried out by the provider; provide necessary information and documentation to competent authorities as required; cooperate with competent authorities upon their reasoned request and take necessary action where required; and comply with registration obligations.

Obligations of importers (Art.23) – Importers must ensure that the AI system is in conformity with the EU AI Act by verifying that the relevant conformity assessment procedure has been carried out by the provider of the high-risk AI system (among other things).

As noted above, an importer is any organisation located or established in the EEA that places on the market an AI system that bears the name or trademark of any organisation outside the EEA. Before placing a high-risk AI system on the market, importers must verify that the conformity assessment procedure has been carried out (Art.43); the necessary technical documentation has been drawn up (Art.11); the AI system bears a CE marking (Art.47); and the provider has appointed an authorised representative (Art.22(1)).

Moreover, importers must refrain from placing high-risk AI systems on the market where they suspect non-compliance or falsification. In fact, where the AI system presents risks to the health, safety, or fundamental rights of persons, importers must inform the provider, authorised representative, and relevant authorities. Importers must indicate their name, registered trade name or trademark, and contact address on the AI system, packaging, or

accompanying documentation. They must also ensure that the storage and transport conditions (that fall under their responsibility) do not compromise compliance with regulatory requirements. Importers also have recordkeeping obligations, as well as obligations regarding cooperation with competent authorities.

Obligations of distributors (Art.24) – Distributors must verify that the AI system bears the required CE marking, is accompanied by a copy of the EU declaration of conformity and instructions for use, and that the provider and the importer of that system (as applicable) have complied with their respective obligations.

As noted above, a distributor is any organisation in the supply chain, other than the provider or the importer, that makes an AI system available on the EEA market. Before making a high-risk AI system available on the market, distributors must verify that the provider and importer of that system have complied with their obligations laid down in Arts.16(b), (c), and 23(3). Like importers, distributors must refrain from placing high-risk AI systems on the market where they suspect non-compliance with the requirements of Arts.8 to 15 (see Chapter 7). Distributors must further ensure that storage and transport conditions (that fall under their responsibility) do not compromise compliance with regulatory requirements. If a distributor considers (or has reason to consider) that a high-risk AI system made available on the market does not conform with the requirements of Arts.8 to 15, they must take necessary corrective actions to bring the AI system into conformity, withdraw it, or recall it. Where the AI system presents risks to the health, safety, or fundamental rights of persons, the distributor must immediately inform the provider or importer of the AI system and the relevant authorities. Distributors also have obligations regarding cooperation with competent authorities.

Responsibilities along the AI value chain (Art.25) –

This provision stipulates the responsibilities of distributors, importers, deployers, and other third parties, considering these parties as providers under certain circumstances.

Any distributor, importer, deployer, or other third party will be regarded as a provider of a high-risk AI system under the EU AI Act (and thus will be subject to the provider's obligations noted above) if they meet any of the following conditions: (i) they put their name or trademark on a high-risk AI system already on the market or in service; (ii) they make a substantial modification to a high-risk AI system already on the market or in service such that it remains a high-risk AI system per Art.6; or (iii) they alter the intended purpose of an AI system, including a GPAI system, that was not initially classified as high-risk, in such a way that it becomes a high-risk AI system pursuant to Art.6.

When these conditions are met, the original provider will no longer be considered the provider for that specific AI system. They must cooperate with the new provider by providing necessary information, technical access, and assistance to fulfil compliance obligations (unless the original provider has specified that their AI system should not be modified into a high-risk AI system).

For high-risk AI systems that are safety components of products covered by EU harmonisation legislation (e.g., in relation to medical devices), the product manufacturer will be considered the provider if the AI system is marketed or put into service under the product manufacturer's name or trademark.

Obligations of deployers (Arts.26 and 27) – Deployers of high-risk AI systems are also subject to specific obligations set out in Art.26. In addition, Art.27 contains specific obligations for deployers of certain high-risk AI systems referred to in Art.6(2) and Annex III (subject to certain exceptions) requiring an assessment of the impact on fundamental rights by such systems.

□ **Obligations of deployers (Art.26) –** Deployers are required to implement appropriate technical and organisational measures to ensure they use high-risk AI systems in accordance with the instructions for use accompanying the AI systems, and assign human oversight to appropriate individuals. These obligations are in addition to any other obligations the deployer may be subject to under EEA, EU, or national law. To the extent that deployers exercise control over the input data, they must ensure such input data is relevant and sufficiently representative considering the intended purpose of the high-risk AI system. Further, deployers must monitor the operation of the AI system, and must inform providers if the AI system poses a risk. Deployers must also keep logs generated by the AI system for an appropriate period (at least six months) and must cooperate with the competent authorities. Deployers that are subject to EU rules on financial institutions are subject to additional obligations – see below.

□ **Obligations for specific types of deployers:**

- **Employers:** Deployers that are employers using a high-risk AI system must inform affected workers and their representatives before using such system.
- **Public authorities:** Deployers that are public authorities or EEA institutions, bodies, offices, or agencies must comply with the registration obligations referred to in Art.49, and inform the provider or distributor if the high-risk AI system is not registered in the EU database.
- **Biometric identification:** Deployers of high-risk AI systems that are used for post-remote biometric identification in the context of criminal investigations must obtain judicial or administrative authorisation in advance or, if immediate deployment is necessary, without undue delay (and in any event within 48 hours). Notably, such systems cannot be used

for an untargeted purpose, or for law enforcement that does not have a direct link to a criminal offence. Each use must be documented in the relevant police file, and also in the relevant annual reports which should be submitted to market surveillance and data protection authorities.

- **Fundamental rights impact assessment for high-risk AI systems (Art.27)** – Prior to deploying a high-risk AI system referred to in Art.6(2) (see Chapter 6), a few categories of deployers must assess the impact of the relevant AI system on the fundamental rights of the affected individuals. The deployers that must do this are: (i) deployers that are governed by public law or are providing public services; and (ii) deployers of high-risk AI systems that are intended to be used to evaluate creditworthiness, carry out risk assessments, and set pricing of life and health insurance (paragraph 5(b) and (c) of Annex III).

For these purposes, deployers are required to perform an assessment that includes: a description of how the high-risk AI system will be used; the period and frequency of use; the categories of people likely to be affected; the specific risks to these individuals or groups; the human oversight measures put in place; and an outline of the risk mitigation measures for handling materialised risks, including governance and complaints mechanisms.

This obligation applies to the first use of the high-risk AI system. Deployers may rely on previously conducted fundamental rights impact assessments or existing impact assessments conducted by the provider in similar cases. Once the assessment is performed, deployers must notify the market surveillance authority of its results, unless they are exempted.

Deployers using AI systems for these purposes (without human review) would also need to consider whether the prohibition on automated decision-making in Art.22 GDPR applies to such AI systems. If any of the requirements in Art.27 EU AI Act have already been fulfilled as a result of completing a Data Protection Impact Assessment under Art.35 GDPR, the deployer is not required to repeat those requirements for the purposes of the EU AI Act.

AI systems intended for use in critical infrastructure are excluded from the requirements of Art.27(1).

It is anticipated that the AI Office will develop a template to assist deployers in meeting these obligations in a simplified manner.

Special rules for providers and deployers that are financial institutions (Arts.17(4), 18(3), 19(2), 26(5) and (6)) – Providers – and in the case of Art.26, deployers – of high-risk AI systems that are subject to EU rules on financial institutions can fulfil their obligations under the EU AI Act (i.e., those relating to management systems, technical documentation, logs, monitoring and logs maintenance) by complying with the requirements under financial services law. In practice, this means that financial institutions and regulated entities must ensure that their EU AI Act compliance program and their financial regulatory compliance measures are synchronised and consistent.

Context and illustrations



Commentary: Managing responsibilities along the AI value chain

The EU AI Act is likely to have a significant impact on commercial contracts relating to AI systems. For high-risk AI systems, managing responsibilities along the AI value chain will involve ensuring that businesses comply with the EU AI Act's requirements. These requirements can be complex, time-consuming, and typically require extensive oversight and communication among the relevant stakeholders. This compliance is unlikely to happen automatically; therefore, providers and deployers are likely to incorporate these obligations into their standard terms, in an attempt to allocate responsibility and potential liability.

It is important that providers, deployers, and other parties involved with high-risk AI systems establish clear contractual obligations amongst various stakeholders and monitor compliance with those obligations. Contracts should include detailed provisions outlining specific allocations of obligations, and incorporate clauses that explicitly specify which party bears responsibility for which aspects of the EU AI Act. Businesses should also consider whether their contracts should contain audit rights that allow them to verify that other relevant actors have fulfilled their respective obligations. When deploying subcontractors, businesses should ensure that they, and other third parties down the supply chain, also comply with the obligations. For example, through flow-down clauses, outlining the consequences of non-compliance and providing mechanisms for mitigation and remediation.



Practical tip: Avoiding duplication

Providers of high-risk AI systems that place the AI system on the market, or put the AI system into service alongside another product (e.g., the software component of a medical device) for which they are subject to obligations regarding a quality management system under relevant sectoral EU law, such as the MDR, may include the quality management aspects listed in Art.17(1) as part of the quality management system for the medical device (software) pursuant to the MDR. In such cases, the provider may be able to limit its obligations under the EU AI Act, to the extent that it can demonstrate that it has already satisfied those obligations under other EU laws.



Commentary: Overlapping regulatory requirements

Certain high-risk AI systems are covered not only by the EU AI Act, but also by EU harmonisation legislation listed in Section A of Annex I to the EU AI Act. This applies, for instance, in the case of AI systems that are intended to be used as a safety component of a medical device.

For such high-risk AI systems:

- The provider should follow the relevant conformity assessment procedure as required under the relevant EU law (e.g., the MDR) and fulfil any of the requirements of Arts.8 to 15 EU AI Act that are not already fulfilled through compliance with that other EU law.
- To ensure consistency, avoid duplication, and minimise additional burdens, providers should consider integrating the necessary elements of a post-market monitoring system (see Art.72(1) to (3)) into their existing compliance monitoring systems, to help ensure that no compliance obligations are missed.

Chapter 09

Notifying authorities and conformity assessment bodies

Executive summary

The EU AI Act requires Member States to establish notifying authorities that are responsible for managing conformity assessment bodies. Notifying authorities inform the Commission and Member States of each conformity assessment body. Notified bodies must meet the required standards of competence and impartiality.

The Commission is responsible for overseeing and ensuring cooperation and coordination among notified bodies.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.28 to 39**, which govern the creation and roles of notifying authorities and conformity assessment bodies.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Notifying authority (Art. 3(19))** – The national authority responsible for setting up and carrying out the necessary procedures for the assessment, designation, and notification of conformity assessment bodies and for their monitoring;
- **Conformity assessment (Art.3(20))** – The process of demonstrating whether the requirements set out in Chapter III, Section 2, relating to a high-risk AI system have been fulfilled;

- **Conformity assessment body (Art.3(21))** – A body that performs third-party conformity assessment activities, including testing, certification, and inspection; and
- **Notified body (Art.3(22))** – A conformity assessment body notified in accordance with the EU AI Act and other relevant EU harmonisation legislation.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Background (Arts.29 – 39) – Under the EU AI Act, a conformity assessment body must apply for notification to be officially recognised as a “notified body”.

Without this notification, a conformity assessment body cannot offer its services to providers for the purposes of conformity assessments (Art.43 – see Chapter 10). In principle, a conformity assessment body that has not yet become a notified body could still provide unregulated services, including consulting services, advice, and training.

Designation and role of notifying authorities (Art.28) – Member States must designate at least one notifying authority. The notifying authority is responsible for assessment, designation, and notification of conformity assessment bodies and for monitoring the performance of conformity assessment bodies.

Notifying authorities must be impartial, avoid conflicts of interest, and ensure that decisions on the notification of conformity assessment bodies (see below) are made by different individuals than those who conduct the assessments of those bodies. To avoid conflicts of interest, notifying authorities cannot provide any of the services that conformity assessment bodies provide, nor can they offer any consultancy services. Notifying authorities must protect the confidentiality of obtained information. Notifying authorities need to have sufficient competent personnel to perform their tasks effectively.

Application process for conformity assessment bodies to become notified bodies (Art.29) – To become a “notified body”, a conformity assessment body must apply for notification to the notifying authority of the Member State in which it is established.

The application must include a description of the conformity assessment activities and the types of AI systems for which the body claims to be competent. If the conformity assessment body has an accreditation certificate issued by a national accreditation body attesting that the conformity assessment body fulfils the requirements laid down in Art.31, the conformity assessment body should include that certificate in its application to the notification authority. If the conformity assessment body does not have an accreditation certificate, it must provide documentary evidence to the notifying authority to demonstrate ongoing compliance with Art.31. Notified bodies already designated under other EU legislation can use relevant documents to support their designation. Once a conformity assessment body has become a notified body, it must keep its documentation updated to allow continuous compliance monitoring.

Notification procedure and requirements (Art.30) – The notification procedure is the procedure by which the national authority notifies the Commission and the Member States that a conformity assessment body has become a notified body.

Using an electronic tool provided by the Commission, national authorities must send a notification to the Commission and the Member States of each conformity assessment body that has satisfied the requirements set out in Art.31. Only a conformity assessment body that has satisfied those requirements can become a notified body.

The notification must include full details of the conformity assessment body, its activities, the types of AI systems it is competent to assess, and an attestation of competence. A conformity assessment body may act as a notified body only if no objections are raised by the Commission or the other Member States within the applicable timeframe for objections (two weeks of a notification when an accreditation certificate is provided, or two months when other documentary evidence is submitted). If objections arise, the Commission must consult with the relevant Member States and the conformity assessment body, after which the Commission must issue a decision on whether the conformity assessment body should be confirmed as a notified body (Art.30).

Operation of notified bodies (Arts.31 – 34) – Notified bodies are required to operate in accordance with mandatory requirements set out in the EU AI Act.

The requirements for notified bodies are laid down in Art.31. In short, a notified body must be a legal entity established under the national law of a Member State. It must have suitable organisational structures, quality management systems, resources, processes, and cybersecurity measures that are suitable for its tasks. A notified body must have a clear organisational structure that ensures confidence in its performance and assessment results. It must be independent from providers of high-risk AI systems and any other parties affected by its assessments and must implement documented procedures to ensure confidentiality of information it receives during assessments. To maintain impartiality, a notified body’s personnel must not be involved in the design, development, marketing, or use of the AI systems it assesses. Similarly to the requirements for notifying authorities (see Art.28), procedures must be in place to safeguard notified bodies.

A conformity assessment body is presumed to comply with the requirements set out in Art.31 if it demonstrates that it conforms to the applicable standards issued by the European Standardisation Organisations (CEN, CENELEC, ETSI) (Art.32).

Where a notified body subcontracts specific tasks connected with the conformity assessment (or uses a subsidiary for that purpose), it must: (i) ensure that the subcontractor (or subsidiary) meets the requirements set out in Art.31; (ii) take full responsibility for the performance of the subcontracted tasks; (iii) obtain the prior agreement of the provider affected by the subcontracting arrangement; and (iv) maintain the relevant records for five years (Art.33).

Operational obligations of notified bodies (Art.34) – A core function of notified bodies is to carry out conformity assessments of high-risk AI systems, in accordance with the procedures set out in Art.43 (see Chapter 10).

Notified bodies should minimise burdens for the providers of high-risk AI systems (especially for micro and small enterprises) while ensuring compliance with required standards. Upon request by the notifying authority, notified bodies must provide all relevant documentation (including the relevant provider's documentation) for assessment and monitoring.

Lists of notified bodies, changes to notifications, and administration of notifications (Arts.35 – 38) – Arts.35 – 38 contain detailed provisions regarding the management and oversight of notified bodies by the Commission.

The core management and oversight roles are as follows:

- The Commission is responsible for assigning an identification number to each notified body and publishing the list of the bodies notified under the EU AI Act (Art.35).
- Each notifying authority must notify the Commission of any relevant changes to the status of a notified body (Art.36).
- The Commission must undertake various administrative functions, including investigating challenges to the competence of notified bodies where necessary, and ensuring coordinating and cooperation between notified bodies (Arts.37 – 38).

Conformity assessment bodies outside the EEA (Art.39) – Conformity assessment bodies that are based in a country outside the EEA may carry out the activities of notified bodies under an agreement between the EU and that country, if they meet the requirements in Art.31, or ensure an equivalent level of compliance.

Context and illustrations



Simplified requirements for notified bodies already designated

For notified bodies which are designated under any other EU harmonisation legislation (e.g., the MDR), documents and certificates linked to those designations may be used to support their designation procedure under the EU AI Act, as appropriate.



Obligation to take out liability insurance

Notified bodies are required to take out appropriate liability insurance for their conformity assessment activities, unless liability is assumed by the Member State in which they are established in accordance with national law or that Member State is itself directly responsible for the conformity assessment (Art.31(9)). This obligation also exists in relation to designated roles under similar EU legislation (including the MDR).



Guidance for requirements to be met by notified bodies

While the EU AI Act provides limited detail regarding the requirements for notified bodies (Art.31), other EU legislation (e.g., the MDR) provides a more extensive catalogue of requirements. Notified bodies under the EU AI Act may be able to look to such requirements for guidance on how the EU AI Act is likely to be interpreted and applied in practice.

Chapter 10

Standards, conformity assessments, certificates, and registration

Executive summary

The EU AI Act creates rules regarding harmonised standards, common specifications, conformity assessments, certificates, and registration of AI systems. High-risk AI systems that conform to harmonised standards or common specifications are, in some circumstances, presumed to comply with the relevant requirements of the EU AI Act. To obtain a conformity assessment, providers must follow assessment procedures and register their high-risk AI systems.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.40 – 49** of the EU AI Act, which generally provide further detail on the processes by which AI systems can actually achieve compliance with relevant technical requirements contained elsewhere in the EU AI Act. Specifically, Arts.40 – 44 and 46 – 49 deal with harmonised standards, common specifications, conformity assessments, certificates, declarations of conformity, CE marking, and registration, while Art.45 deals with information obligations on notified bodies relating to certificates.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Harmonised standard** (Art.3(27)) – A standard adopted on the basis of a request made by the Commission for the application of EU harmonisation legislation (as defined in Art.2(1)(c) of Regulation (EU) 1025/2012).
- **Common specification** (Art.3(28)) – A set of technical specifications that prescribes technical requirements to be fulfilled by a product, process, service, or system, and which lays down one or more of requirements in Art.2(4)(a) – (d) of Regulation (EU) 1025/2012.
- **CE marking** (Art.3(24)) – A marking by which a provider indicates that an AI system is in conformity with the requirements set out in Arts.8 – 15, and other applicable EU harmonisation legislation providing for its affixing.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Simplifying compliance for businesses: Harmonised standards, common specifications, and a presumption of conformity – High-risk AI systems and GPAI models conforming with certain harmonised standards or common specifications are presumed to be in conformity with certain requirements set out in the EU AI Act (Art.40). This presumption may help to simplify compliance processes for businesses.

The EU AI Act establishes detailed requirements for high-risk AI systems and providers of GPAI models (see Arts.8 – 15 and 53 – 55, and Chapters 7 and 13). Arts.40 and 41 provide a simpler route to compliance: High-risk AI systems and GPAI models conforming to harmonised standards published in the [Official Journal](#), or common specifications provided by the Commission, will benefit from **presumed** conformity with the aforementioned requirements.

These harmonised standards are required to cover all relevant requirements and obligations under the EU AI Act, and provide further detail for reporting deliverables and documentation processes. In other words, harmonised standards will (in principle) provide a route through which businesses can achieve compliance with the relevant requirements of EU AI Act without having to conduct their own independent analysis of how to achieve such compliance. This may be advantageous to many businesses – especially where the interpretation of the EU AI Act’s requirements is unclear.

The interplay between harmonised standards and common specifications (Arts.40 and 41):

- **Harmonised standards** – The Commission will ask existing EU standardisation organisations to develop harmonised standards. Some standardisation organisations have already been tasked with developing harmonised standards for high-risk AI systems. These are expected to be published by the end of 2025; however, at the time of writing, no harmonised standards have been published.
- **Common specifications** – If: (i) a standardisation organisation has been asked to draft harmonised standards; and (ii) the relevant standardisation organisation fails to accept the request, misses the relevant deadline, does not sufficiently address fundamental rights concerns, or develops standards which do not comply with the request, then the Commission may adopt implementing acts establishing common specifications for AI requirements and obligations (Art.41).

This is intended as a backstop against absent or insufficient harmonised standards: The Commission's common specifications will be repealed if harmonised standards are later published in the Official Journal (Art.41(4)).

Where providers of high-risk AI systems/GPAI models do not comply with the common specifications, they are required to justify that non-compliance (e.g., by demonstrating that they have implemented suitable alternative means of achieving compliance) (Art.41(5)). Member States can also challenge a common specification if they consider that it does not fully meet the requirements elsewhere in the EU AI Act, prompting review and amendment of the specifications, if necessary (Art.41(6)).

High-risk AI systems may benefit from a presumption of conformity with certain requirements even if there are no harmonised standards or common specifications – High-risk AI systems will be presumed to comply with certain requirements under the EU AI Act, under two specific conditions (Art.42):

- **Data context compliance** – If high-risk AI systems are trained and tested using data that reflects the specific geographical area, behaviour, context, or function they will be used for, they are presumed to comply with the relevant requirements in Art.10(4) of the EU AI Act (Art.42(1)).
- **Cybersecurity compliance** – If high-risk AI systems have been certified or have a statement of conformity issued under a cybersecurity scheme pursuant to the EU Cybersecurity Act, with references published in the Official Journal, they are presumed to comply with the cybersecurity requirements set out in Art.15 of the EU AI Act, insofar as the cybersecurity certificate or statement of conformity covers those requirements (Art.42(2)).

Conformity assessment and certification (Arts.43 – 44) – Providers of high-risk AI systems will still need to undergo conformity assessment procedures and certification (Arts.43 – 44). Limited exceptions to conformity assessment procedures are discussed below (Art.46).

The conformity assessment procedure options laid down in Art.43 vary according to the specific high-risk AI system involved:

- **High-risk AI systems listed in the area of biometrics** – Providers of high-risk AI systems listed in paragraph 1 of Annex III (i.e., biometric systems) that have applied the harmonised standards or common specifications (Arts.40 and 41) can choose between the following options:
 - **Conformity assessment procedure based on internal controls** (Annex VI) – In this procedure the provider of the high-risk AI system must verify that the quality management system, technical documentation, and development processes all align with the relevant regulatory requirements under the EU AI Act (Art.43(1)(a)).
 - **Conformity assessment procedure based on quality management system and technical documentation assessment by a notified body** (Annex VII) – The notified body must assess whether the provider satisfies the requirements set out in detail in Annex VII, regarding the quality management system and technical documentation (Art.43(1)(b)) (see Chapter 9 for an explanation of notified bodies). The notified body may conduct periodic audits and additional tests to ensure ongoing compliance.

- The provider of a high-risk AI system must follow the procedure set out in Annex VII (and involve a notified body) if any of the following apply (Art.43(1)):
 - No harmonised standards or common specifications are available.
 - The provider has not fully applied the harmonised standards or common specifications.
 - Existing harmonised standards include restrictions (and the Annex VII procedure must be applied only to the part of the standard that was restricted).
- **Choice of notified bodies** – For the purposes of the conformity assessment procedure under Annex VII, providers can choose any available notified body (except in relation to AI systems used by law enforcement, immigration, or asylum authorities or by EU institutions, which require assessment by the market surveillance authority) (Art.43(1)).
- **Internal controls for certain high-risk AI systems**
 - Providers of high-risk AI systems listed in paragraphs 2 – 8 of Annex III (i.e., AI systems relating to critical infrastructure, education, employment, essential services, law enforcement, immigration, justice, and democracy) must follow the conformity assessment procedure based on internal controls (Annex VI) which does require the involvement of a notified body (Art.43(2)).
- **Other EU harmonisation legislation** – Providers must follow the relevant conformity assessment procedures for high-risk AI systems covered by other EU harmonisation laws, such as the MDR, incorporating the EU AI Act's requirements (Art.43(3)).
- **Substantial modifications** – High-risk AI systems must undergo a new conformity assessment if they are substantially modified. For high-risk AI systems that continue to learn after being placed on the market or put into service, changes to the AI system and its performance that have been pre-determined by the provider before the initial conformity assessment, and that are part of the AI system's technical documentation, do not constitute a substantial modification and do not require a new conformity assessment (Art.43(4)).

- **The Commission's authority** – The Commission can update Annexes VI and VII and amend conformity assessment requirements based on technical progress and the effectiveness of internal controls (Art.43(5)).

Certificates issued by notified bodies in accordance with Annex VII must be in a language easily understood by relevant authorities in the Member State in which the notified body is established (Art.44(1)). Certificates are valid for up to five years for AI systems that fall under Annex I (i.e., AI systems falling within Art.6(1)(a) or (b), relating to product safety), and four years for Annex III AI systems (i.e., high-risk AI systems under Art.6(2)) with possible extensions upon re-assessment (Art.44(2)). If an AI system no longer meets the requirements, the notified body may suspend, withdraw, or restrict the certificate unless corrective actions are taken (Art.44(3)).

Notified bodies that issue, refuse, restrict, suspend, or withdraw certificates, supplements to certificates, and quality management systems approvals in accordance with Annex VII must share the required information with the relevant notifying authority and other notified bodies (in some instances, the notified body must provide this information proactively; in other instances, notified bodies are only required to provide it on request). Notified bodies must safeguard the confidentiality of information that they obtain in accordance with Art.78 (Arts.45(1)-(4)) (see Chapter 9 for an explanation of notifying authority).

Exceptions to the conformity assessment procedure (Art.46) – There are limited (exceptional) circumstances which allow for a derogation from the usual conformity assessment procedure:

- **Limited approvals in exceptional circumstances**
 - In exceptional circumstances, a market surveillance authority may authorise specific high-risk AI systems to be placed on the market or put into service without a conformity assessment (e.g., for public security, the protection of life and health of persons, environmental protection, or the protection of key industrial and infrastructure assets) provided that the authority concludes that the high-risk AI system complies with the requirements of the EU AI Act (see Chapter 7) (Art.46(1)). This authorisation is temporary, and the necessary conformity assessment procedures still need to be carried out without undue delay (Arts.46(1) and (3)).

- **Limited exemption for public security emergencies or threats to safety** – Law enforcement or civil protection authorities may temporarily put a specific high-risk AI system into service without authorisation from a market surveillance authority where necessary due to: (i) urgency for exceptional reasons of public security; or (ii) a specific, substantial, and imminent threat to life or physical safety, provided that such authorisation is requested during or after the use without undue delay (Art.46(2)). If the market surveillance authority then refuses authorisation, the use of the specific high-risk AI system must cease with immediate effect, and the results and outputs of such use must be immediately discarded (Art.46(2)).

Where such exceptional circumstances apply, the relevant market surveillance authority must notify the Commission and other Member States. The market surveillance authority is not required to disclose sensitive operational data in relation to the activities of law-enforcement authorities (Art.46(3)). Member States and the Commission then have 15 calendar days to object to the authorisation, which may trigger the need for further consultations (Art.46(5)). Where the Commission ultimately decides that the authorisation is unjustified, the relevant market surveillance authority shall withdraw its authorisation (Art.46(6)).

Declaration of conformity (Art.47) – Providers of high-risk AI systems must create and maintain a written, machine readable, signed declaration of conformity for each high-risk AI system. That declaration of conformity must be kept, and made available to national competent authorities on request, for ten years after the AI system is placed on the market or put into service.

- **Content** – Each declaration of conformity must:
 - (i) identify the high-risk AI system for which it has been created;
 - (ii) state that the relevant high-risk AI system meets the requirements of Arts.8 – 15 (see Chapter 7);
 - (iii) contain the information set out in Annex V (which lists categories of information that a declaration of conformity must contain); and
 - (iv) be translated into a language that can be easily understood by the national competent authorities of the Member States in which the high-risk AI system is being placed on the market or put into service (Arts.47(1)-(2)). For providers whose AI systems are available across the EU, this may result in the need for a significant number of translated versions of the declaration of conformity.
- **Maintenance** – A declaration of conformity must:
 - (i) be made available to the relevant national competent authorities for ten years after the high-risk AI system has been placed on the market or put into service; and
 - (ii) be kept up-to-date by the provider. It is important to note that the Commission is able to amend the contents of Annex V, which may mean that existing declarations of conformity have to be updated when such amendments take place (Arts.47(1), (4), and (5)).
- **Overlapping legislation** – If a high-risk AI system falls under other EU harmonisation legislation requiring a declaration of conformity, a single declaration should be created to cover each piece of applicable legislation, and that declaration should clearly explain what legislation it covers (Art.47(3)).
- **Assumption of responsibility** – By creating the declaration of conformity, the provider assumes responsibility for compliance with the requirements of Arts.8 – 15 (see Chapter 7) (Art.47(4)).

Marking (Art.48) – CE marking signifies that a product meets certain health, safety, and environmental requirements of EU law. High-risk AI systems must be affixed with: (i) a CE marking that is visible, legible, and indelible; and (ii) the identification number of the notified body, where applicable:

- **CE marking** – CE marking is required for high-risk AI systems (Art.16(h)). High-risk AI systems provided within physical products need a physical CE marking (which can be complemented with a digital one). High-risk AI systems provided digitally need digital CE marking, which must be easily accessible via the AI systems interface, machine-readable code, or other electronic means (Rec.129; Art.48(2)). Alternatively, where digital CE marking is not possible or not appropriate due to the nature of the high-risk AI system, CE marking must be affixed to the packaging or accompanying documentation, as appropriate (Art.48(3)). Failure to affix CE marking, or incorrect use of CE marking, may lead to enforcement by the relevant market surveillance authorities (Art.83(1)(a)-(b)).
- **Notified body identification** – Where applicable, the CE marking must be followed by the identification number of the notified body responsible for the conformity assessment procedures set out in Art.43. The identification number must be included in any promotional material referring to the high-risk AI system's fulfilment of the requirements for CE marking (Art.48(4)).

CE marking is subject to the general principles set out in Article 30 of Regulation (EC) 765/2008. If a high-risk AI system is also subject to other EU laws that require CE marking, the CE marking on the high-risk AI system must clarify that the high-risk AI system also fulfils the requirements of those laws (Arts.48(1) and (5)).

Registration (Art.49) – Certain AI systems come with registration obligations:

- **Most of the high-risk AI systems listed in Annex III** – The provider or authorised representative, and/or certain deployers of AI systems listed in Annex III, must register themselves and their system in the EU database before placing the AI system on the market or putting it into service (Arts.49(1) and (3)). This applies to all high-risk AI systems listed in Annex III, except critical infrastructure systems, which must be registered at national level (Art.49(5)).
- **High-risk AI systems in the areas of biometrics, law enforcement, migration, asylum, and border control** – Such high-risk AI systems must be registered in a secure, non-public section of the database, with access limited to the Commission and certain national authorities (Art.49(4)).
- **AI systems that are not deemed high-risk by the provider** – Under Art.6(3), a provider of an AI system listed in Annex III can determine that the AI system is not high-risk (see Chapter 6). The provider or authorised representative must register themselves and the relevant AI system in the EU database before placing that AI system on the market or putting it into service (Art.49(2)).

Context and illustrations



Stay up-to-date: Commission implementing decision on a standardisation request to CEN/CENELEC

By implementing decision of 22 May 2023, the Commission has mandated the European standardisation organisations CEN/CENELEC to develop technical standards regarding the mandatory requirements for high-risk AI systems by 30 April 2025. However, at the time of writing, it appears this deadline has not been met, and the standards remain outstanding. The Commission will then decide whether the standards developed by CEN/CENELEC comply with the standardisation mandate. If this is the case, a reference to the standard concerned will be published in the Official Journal. The CEN/CENELEC Joint Technical Committee 21 is in the process of adapting standards from ISO/IEC and developing new standards. The respective work programme of the Committee together with the current status of drafting and approval as well as the forecasted timeline is available [here](#).



Example: Navigating high-risk AI system compliance

Company X provides a high-risk AI system in the area of education (these fall within paragraph 3 of Annex III). The high-risk AI system will therefore need to comply with the requirements set out in Arts.8 – 15 of the EU AI Act.

If harmonised standards or common specifications have been developed and published in the Official Journal for Company X's high-risk AI system, and Company X's high-risk AI system complies with those standards or specifications, it will benefit from a **presumption** of conformity (Arts.40 and 41).

In order to benefit from this presumption, Company X will need to follow the conformity assessment procedure based on internal controls (as referred to in Art.43(2); Annex VI).

Company X's AI system will not need a certificate issued by a notified body, because its conformity with EU AI Act requirements is not being assessed by a notified body (Art.44).

Company X must register itself and the relevant AI system in the EU database before placing the AI system on the market or putting it into service (Art.49(1)).



Commentary: Challenges for providers in view of conformity assessments and certification processes

Providers may face several challenges regarding adherence to the conformity assessments. Specifically, adhering to the various conformity assessment procedures, including internal control and reviews by the chosen notified body, may be complex and resource-intensive. These processes can be time-consuming and involve additional administrative tasks. Since conformity certificates only have a shelf-life of four to five years, it is essential that providers introduce ongoing compliance monitoring practices to avoid the risk of a certificate being revoked due to non-compliance. These challenges necessitate significant investments in compliance infrastructure and organisational transitions.

Therefore, providers should: (i) be aware of the different conformity assessments and choose the appropriate option according to their specific high-risk AI systems; (ii) be aware of templates for required documentation, such as quality management systems, technical documentation, and development process details; and (iii) regarding the certification process, providers should closely analyse the processes for obtaining, renewing, and maintaining conformity certificates, including the relevant timelines to effectively navigate the complexities of the conformity assessment and certification processes.

Chapter 11

Transparency obligations for certain AI systems

Executive summary

Article 50 of the EU AI Act contains transparency obligations that apply to specific types of AI systems.

For these AI systems, providers or deployers are required to inform users that they are interacting with an AI system or AI-generated output (as opposed to a human or human-generated output), subject to some exceptions.

These obligations exist alongside additional transparency obligations in the EU AI Act that apply to high-risk AI systems or other categories, which are detailed in Chapters 7 (in relation to high-risk AI systems) and 13 (in relation to GPAI models).

Relevant sections of the EU AI Act

This Chapter focuses on **Art.50** of the EU AI Act – specifically, transparency obligations on providers and deployers of certain AI systems and applicable exceptions. This Chapter also includes insights on the relevant definitions in Art.3, to the extent that those definitions relate to the scope of the EU AI Act.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **AI system (Art.3(1))** – A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.
- **Deep fake (Art.3(60))** – AI-generated or manipulated image, audio, or video content that resembles existing persons, objects, places, entities, or events and would falsely appear to a person to be authentic or truthful.
- **Transparency (Rec.27)** – The term is not explicitly defined in the EU AI Act. However, Recital 27 explains that the principle of transparency means that AI systems are developed and used in a way that allows appropriate traceability and explainability, while making humans aware that they communicate or interact with an AI system, as well as duly informing deployers of the capabilities and limitations of that AI system and affected persons about their rights.
- **Emotion recognition system (Art.3(39))** – An AI system for the purpose of identifying or inferring emotions or intentions of natural persons on the basis of their biometric data.
- **Biometric categorisation system (Art.3(40))** – An AI system for the purpose of assigning natural persons to specific categories on the basis of their biometric data, unless it is ancillary to another commercial service and strictly necessary for objective technical reasons.

A full list of defined terms can be found in the [Glossary](#).

Analysis

In-scope AI systems – Art.50 of the EU AI Act imposes information and notification obligations on providers or deployers of certain types of AI systems. Those AI systems and the applicable obligations are as follows:

- **AI systems interacting with individuals** (Rec.132; Art. 50(1)) – Providers must ensure that AI systems that interact directly with individuals are designed in a way that informs them that they are interacting with an AI system (i.e., the provider needs to notify users that they are interacting with an AI system). This category could include, for example, chatbots, voice-assistants, and robo-services, but would not include AI systems interacting with other AI systems.
- **AI systems generating synthetic content** (Rec.133; Art. 50(2)) – Providers of AI systems that generate synthetic audio, image, video, or text content must ensure that these outputs are marked in a machine-readable format and detectable as artificially generated or manipulated, by using effective technical solutions as far as feasible (such as watermarking the content or using metadata identification measures). This obligation applies to a broad scope of AI systems, including text generators, chatbots, audio generating systems, image editing software, video editing tools, and so on, where those systems include an AI component that is used to generate any output.
- **AI systems involving emotion recognition or biometric categorisation** (Rec.132; Art. 50(3)) – Deployers of AI systems that are used for the purposes of emotion recognition system or biometric categorisation must: (i) inform affected individuals that those AI systems are in use; and (ii) process the personal data of those individuals in accordance with applicable data protection laws (e.g., the GDPR). These obligations will apply to deployers of AI systems that, for example, track viewers' emotional reactions to ads, or assign individuals to specific categories relating to factors such as height, biomechanics, eye colour, personal preferences, etc.

- **AI systems generating/manipulating deep fakes** (Rec.134; Art.50(4)) – Deployers of an AI system that generates or manipulates image, audio, or video content constituting a deep fake must disclose that the content has been artificially generated or manipulated. Where the deep fake forms part of an artistic, creative, satirical, or similar work, the disclosure is only required to occur in an appropriate way without hampering the enjoyment of that work.
- **AI systems generating/manipulating text** (Rec.134; Art.50(4)) – Deployers of an AI system that generates or manipulates text, which is published to inform the public on matters of public interest (e.g., news publications, safety warnings, health alerts, information affecting fundamental rights, etc.), must disclose that the text has been artificially generated or manipulated.

These obligations are further detailed in the table below.

In the absence of further guidance, it appears that these obligations have a cumulative effect. For example, a deep fake generated using user prompts will likely trigger both a disclosure obligation (Art.50(4)) and a marking obligation to indicate that the deep fake has been artificially generated (Art.50(2)).

As with other parts of the EU AI Act, these obligations are stated to exist in parallel and without prejudice to additional transparency obligations in the EU AI Act that apply to high-risk AI systems and other transparency obligations under EU or national laws (Art.50(6)). Accordingly, these obligations also appear to apply cumulatively. See Chapters 7 (in relation to high-risk AI systems) and 13 (in relation to GPAI models) for more details.

Therefore, businesses will need to carefully review their AI systems to ensure that each of the applicable transparency obligations across the EU AI Act is satisfied in respect of each such system.

Exceptions – Art.50 of the EU AI Act provides certain exceptions, where the information and notification obligations outlined above are inapplicable. Those exceptions are as follows:

- **For all AI systems listed above** – The obligations listed above do not apply to an AI system if that AI system: (a) falls outside of the scope of the EU AI Act under Art.2 (e.g., any individuals using AI systems in the course of a purely personal non-professional activity); or (b) is authorised by law for law enforcement purposes (in some cases, subject to appropriate safeguards to rights and freedoms of third parties). However, exemption (b) does not apply where an AI system is available for use by the public for the purpose of reporting a criminal offence (e.g., a chatbot on a law enforcement authority’s website).
- **AI systems interacting with individuals** – The notification obligations in Art.50(1) do not apply if it is obvious to reasonably well-informed, observant, and circumspect individuals that they are interacting with an AI system, taking into account the circumstances and context. Businesses should be cautious when relying on this exemption, as a court or regulator may narrowly construe the level of information that a “reasonably well-informed” user might be deemed to have.
- **AI systems generating synthetic content** – The obligations in Art.50(2) do not apply to AI systems that: (a) are used primarily for standard editing assistance; or (b) do not substantially alter the input data provided by deployers. Neither the EU AI Act nor its Recitals provide clarity on what type of AI tools were intended to be caught by this exception; however, there may be further clarity regarding this exception once regulatory guidance is published.
- **AI systems generating/manipulating text** – The obligations in Art.50(4) do not apply to AI-generated content where: (a) the content has undergone a process of human review or editorial control; and (b) an individual holds editorial responsibility for the publication of the content.

Form of notification requirements (Rec.132; Art.50(5))

– For each of the specific AI systems outlined above, the EU AI Act imposes certain requirements on the form and timing of the required notifications, requiring providers or deployers to provide the required information:

- In a clear and distinguishable manner.
- Taking into account the characteristics of vulnerable users, where appropriate.
- At the latest, at the time of the first interaction or exposure to the AI system by individuals (which appears akin to point-in-time notice requirements under the GDPR).
- In a manner that conforms to applicable accessibility requirements.

Future codes of practice and guidelines (Rec.135;

Arts.50(7) and 96(1)(d)) – To try to promote consistent application of the Art.50 transparency obligations, the EU AI Act provides multiple means for further guidance on the practical implementation of the obligations under this Art.50:

- The AI Office is responsible for facilitating the creation of codes of practice at an EU level, in relation to the detection and labelling of artificially generated or manipulated content (i.e., Art.50(2) and (4)), which the Commission can approve.
- The Commission has been tasked with developing guidelines on the practical implementation of the Art.50 transparency obligations (Art.96(1)(d)) and, if the codes of practice referred to above are considered inadequate, can adopt an implementing act of common rules for implementing the obligations.
- The AI Office has stated on its website that it will [issue further guidance](#) for providers and deployers on the obligations in Art.50 (which becomes applicable on 2 August 2026).

Therefore, businesses can expect to receive more clarity in the future on the application of these transparency obligations. Until then, it will be prudent to take a cautious and comprehensive approach to compliance, particularly if a business intends to rely on an exception.

In-scope AI systems under Art.50 of the EU AI Act

The table below summarises the transparency obligations that apply to providers and deployers of various types of AI systems. As with other parts of the EU AI Act, these obligations are stated to exist in parallel and without prejudice to additional transparency obligations in the EU AI Act that apply to high-risk AI systems and other transparency obligations under EU or national laws (Art.50(6)). Accordingly, these obligations also appear to apply cumulatively. See Chapters 7 (in relation to high-risk AI systems) and 13 (in relation to GPAI models) for more details.

In each case, the required information must be presented clearly and distinctly, taking into account the needs of vulnerable users, and must be provided no later than the first time that individuals interact with, or are exposed to, the AI system (Art.50(5)).

Type of AI system	Description of obligation	Examples	Exceptions
AI systems intended to interact directly with individuals (Rec.132; Art. 50(1))	<u>Providers</u> must ensure that AI systems that are intended to interact directly with individuals are designed in a way that means affected individuals are informed that they are interacting with an AI system (i.e., the provider needs to make it clear to individuals that they are interacting with an AI system).	This category includes AI systems that are intended to interact directly with individuals – e.g., chatbots, voice-assistants, and robo-services. It does not include AI systems designed to interact exclusively with other AI systems or other non-human systems.	The obligations do not apply: <ul style="list-style-type: none"> □ If the AI system falls outside of the scope of the EU AI Act under Art.2 (e.g., any individuals using AI systems in the course of a purely personal non-professional activity). □ If the AI system is authorised by law for law enforcement purposes, except where that AI system is available for use by the public for the purpose of reporting a criminal offence (Art.50(1)). □ If it would be obvious to reasonably well-informed, observant, and circumspect individuals that they are interacting with an AI system, taking into account the circumstances and context (Art.50(1)).
AI systems generating synthetic content (Rec.133; Art. 50(2))	<u>Providers</u> of AI systems that generate synthetic outputs in audio, image, video, or text formats must ensure that these outputs are marked in a machine-readable format that identifies them as artificially generated or manipulated, by using effective technical solutions as far as possible (such as watermarking the content or using metadata identification measures).	This obligation applies to a broad scope of AI systems including text generators, chatbots, audio generating systems, image editing software, video editing tools, and so on, where those systems include an AI component that is used to generate any output.	The obligations do not apply: <ul style="list-style-type: none"> □ If the AI system falls outside of the scope of the EU AI Act under Art.2 (e.g., any individuals using AI systems in the course of a purely personal non-professional activity). □ If the AI system is authorised by law for law enforcement purposes, except where that AI system is available for use by the public for the purpose of reporting a criminal offence (Art.50(2)). □ If the AI system is used primarily for standard editing assistance (Art.50(2)). □ If the AI system does not substantially alter the input data provided by deployers (Art.50(2)).

Type of AI system	Description of obligation	Examples	Exceptions
AI systems involving emotion recognition or biometric categorisation (Rec.132; Art. 50(3))	<u>Deployers</u> of AI systems that are used for the purposes of emotion recognition or biometric categorisation must: (i) inform affected individuals that those AI systems are in use; and (ii) process the personal data of those individuals in accordance with applicable data protection laws (e.g., the GDPR). In this context, “emotion” includes states such as happiness, sadness, anger, surprise, disgust, embarrassment, excitement, shame, contempt, satisfaction, and amusement. It does not include physical states, such as pain or fatigue (Rec.18).	These obligations will apply to deployers of AI systems that, for example, track viewers’ emotional reactions to ads, or assign individuals to specific categories relating to factors such as height, biomechanics, eye colour, personal preferences, etc.	The obligations do not apply: <ul style="list-style-type: none"> □ If the AI system falls outside of the scope of the EU AI Act under Art.2 (e.g., any individuals using AI systems in the course of a purely personal non-professional activity); or □ If the AI system is authorised by law for law enforcement purposes, except where that AI system is available for use by the public for the purpose of reporting a criminal offence (Art.50(3)).
AI systems generating/manipulating images and videos, or creating “deep fakes” (Rec.134; Art.50(4))	<u>Deployers</u> of an AI system that generates or manipulates images, audio, or video content constituting a deep fake must disclose that the content has been artificially generated or manipulated. Where the deep fake forms part of an artistic, creative, satirical, fictional, or similar work, the disclosure is only required to occur in an appropriate way without hampering the enjoyment of that work.	This category includes all “deep fakes” and any AI-generated or manipulated image, audio, or video content that resembles existing persons, objects, places, entities, or events and would falsely appear to a person to be authentic or truthful.	The obligations do not apply: <ul style="list-style-type: none"> □ If the AI system falls outside of the scope of the EU AI Act under Art.2 (e.g., any individuals using AI systems in the course of a purely personal non-professional activity). □ If the AI system is authorised by law for law enforcement purposes, except where that AI system is available for use by the public for the purpose of reporting a criminal offence (Art.50(4)).
AI systems generating/manipulating text or creating “fake news” (Rec.134; Art.50(4))	<u>Deployers</u> of an AI system that generates or manipulates text, which is published to inform the public on matters of public interest (e.g., news publications, safety warnings, health alerts, information affecting fundamental rights, etc.) or creating “fake news” must disclose that the text has been artificially generated or manipulated.	This category could include, for example, AI systems that are used to produce text-based outputs for the purposes of creating “fake news”.	The obligations do not apply: <ul style="list-style-type: none"> □ If the AI system falls outside of the scope of the EU AI Act under Art.2 (e.g., any individuals using AI systems in the course of a purely personal non-professional activity). □ If the AI system is authorised by law for law enforcement purposes, except where that AI system is available for use by the public for the purpose of reporting a criminal offence (Art.50(4)). □ To AI-generated content where: (a) the content has undergone a process of human review or editorial control; and (b) an individual holds editorial responsibility for the publication of the content (Art. 50(4)).

Context and illustrations



Commentary: The broad and somewhat vague scope of Article 50

A number of the concepts addressed in Art.50, such as deep fakes and AI-generated outputs, largely remain undefined or unspecific, creating uncertainty in the interpretation and implementation of Art.50.

One area of uncertainty arises in the determination of whether the obligations in Art.50 apply to a given AI system. For example, for AI systems that generate synthetic content, a literal reading of Art.50(2) would indicate that the obligation only applies to providers of AI systems that “generate” the content. While this would capture “new” synthetic content, it is unclear if it would also capture slight modifications that are made to existing content.

Another area of uncertainty arises in the practical (and technical) application of the obligations and, in turn, the threshold for a provider or deployer to be deemed compliant. For example, technical solutions used for the marking obligations under Art.50(2) are required to be “interoperable” (among other things) as far as technically feasible. However, it is often not possible for watermarks (which are commonly applied across various media, such as text and videos) to be accurately and uniformly detected – creating an interoperability issue.

In the absence of further guidance from the AI Office or the Commission, businesses are likely to face uncertainty around exactly which AI systems (and their output) are in scope, and how these transparency obligations will apply in practice. Businesses should, therefore, take a cautious approach to compliance with Art.50.



Commentary: The interplay between transparency obligations under the EU AI Act and the DSA

There is some overlap between Art.50 of the EU AI Act and the DSA, particularly in relation to the identification and mitigation of systemic risks. Any organisation falling within the scope of both laws should consider the potential overlap between obligations.

Art.2 clarifies that the EU AI Act is intended to be “without prejudice to” provisions in other EU regulations such as the DSA. Therefore, obligations under Art.50 will operate in parallel with obligations under the DSA.

Art.35 DSA requires providers of VLOPs and VLOSEs to implement reasonable, proportionate, and effective mitigation measures tailored to the specific systemic risks identified. It lists potential mitigation measures that providers should consider implementing in order to meet this obligation (e.g., the prominent marking of deep fake content). The wording of this is similar to the definition of “deep fake” under the EU AI Act.

The Commission has issued guidelines under the DSA for VLOPs and VLOSEs to address systemic electoral risks. These require generative AI content to be detectable (e.g., via watermarks), and that deep fake content is clearly labelled or otherwise distinguishable through prominent markings. These obligations mirror the transparency obligations contained in Art.50 of the EU AI Act.

Deep fake content which has not been marked in compliance with Art.50(4) of the EU AI Act may also constitute “illegal content” under the DSA (see Art.3(h) DSA).



Commentary: The overlap between transparency obligations under the EU AI Act and the GDPR

Businesses face the challenge of complying with parallel transparency obligations under the EU AI Act and under Arts.13 and 14 GDPR. The EU AI Act is stated to be “without prejudice to” the GDPR, and therefore (in principle) does not affect obligations arising from it.

To comply with GDPR transparency obligations, businesses will typically publish: (i) an external privacy policy on their website for third parties; and (ii) an internal privacy notice on the intranet (or employee handbook) for staff. In addition, businesses will need to consider how to provide suitable AI transparency information to third parties and to their own staff (noting that, as set out above in this Chapter, the EU AI Act requires transparency information to be provided either together with the AI system, or on the output of the AI system, meaning that a single “AI policy” will likely not be sufficient in most cases).

Given that most international businesses are subject to multiple transparency obligations under different laws around the world, some businesses are likely to include AI transparency information in their privacy notices where possible. However, it is essential for businesses to ensure that information provided in their AI transparency notices is consistent with the information provided in their privacy notices. This may be a challenge as these documents will likely need to be updated over time. There is no perfect solution to these challenges. Each business will need to consider the available options and implement the solution that best fits their needs.

Chapter 12

GPAI models – Classification rules

Executive summary

Under the EU AI Act, certain GPAI models are deemed to have “*systemic risk*” if they have “*high-impact capabilities*”. The EU AI Act sets out the circumstances in which GPAI models will be presumed to have such capabilities.

In addition, even where a GPAI model does not have “*high-impact capabilities*”, the Commission has wide powers to decide that such a GPAI model has “*systemic risk*” based on other criteria.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.51 – 52** of the EU AI Act – specifically: (i) how providers can determine whether a GPAI model has high-impact capabilities (and therefore has systemic risk); (ii) the procedure providers must use to notify the Commission of such a determination; and (iii) the criteria the Commission can use to designate a GPAI model as having systemic risk (even where a provider has not notified it that a GPAI model has high-impact capabilities).

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model, or that has an AI system or a GPAI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
 - **Systemic risk** (Art.3(65)) – A risk that is specific to the high-impact capabilities of GPAI models, having a significant impact on the EU market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain.
 - **GPAI system** (Art.3(66)) – An AI system which is based on a GPAI model and which has the capability to serve a variety of purposes, both for direct use as well as for integration in other AI systems.
 - **GPAI model** (Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.
- A full list of defined terms can be found in the [Glossary](#).

Analysis

Assessing which GPAI models have systemic risk – Providers should assess whether their GPAI models have “high-impact capabilities”, but the Commission can also take action to designate GPAI models. The EU AI Act imposes additional compliance obligations on GPAI models that have systemic risk. Under Art.51(1) of the EU AI Act, a GPAI model will be deemed to have systemic risk in two circumstances:

First, a GPAI system has systemic risk if it has “*high-impact capabilities*” (Art.51(1)(a)), in which case the provider must notify the Commission (Art.52(1)). “*High-impact capabilities*” are described as “*capabilities that match or exceed the capabilities recorded in most advanced general-purpose AI models*” (Rec.111). This indicates that providers need to compare their GPAI models against the capabilities of other advanced GPAI models (to the extent that this information is available to them). Art.51(1)(a) makes specific reference to the capabilities of GPAI models being assessed on the basis of “*appropriate technical tools and methodologies, including indicators and benchmarks*”. The predominant methodology set out in the EU AI Act to evaluate GPAI model capabilities is the cumulative amount of computation used for training – if this figure exceeds 10²⁵, measured in floating point operations (FLOPs), the GPAI model will be presumed to have high-impact capabilities (Rec.111 and Art.51(2)).

There is no further elaboration on what criteria such indicators and benchmarks should consider; however, the Commission has the ability to pass delegated acts to supplement the indicators and benchmarks referred to. The threshold of 10²⁵ FLOPs can also be adjusted to reflect the state of the art, such as improvements in hardware efficiency or algorithmic improvements (Art.51(3)).

The EU AI Act does not directly address the inherent information gap that is likely to arise from the fact that many GPAI models are likely to be proprietary, meaning that the exact extent of their capabilities cannot necessarily be known or assessed by other providers.

Second, Art.51(1)(b) gives the Commission the power to decide (either “ex officio” or based on input from the Scientific Panel – see Chapter 17) that a GPAI model has equivalent capabilities or impact to those set out in Art.51(1)(a). Under Art.90(1), the Scientific Panel can provide an alert to the Commission where it has reason to suspect: (i) that a GPAI model poses “concrete identifiable risk at EEA level”; or (ii) that a GPAI model meets the conditions of Art.51(1)(a).

In either case, the Commission will assess the capabilities of the GPAI model independently with regard to the criteria set out in Annex XIII (Arts.51(1)(b) and 52(4)). These criteria include:

- The number of parameters of the GPAI model.
- The quality or size of the data set (e.g., measured through tokens).
- The amount of computation used for training the model, measured in FLOPs or indicated by a combination of other variables such as the estimated cost of training, the estimated time required for training purposes, or the estimated energy consumption for training purposes.

- The input and output formats of the model, such as text to text (large language models), text to image, multi-format output, and the state-of-the-art thresholds for measuring high-impact capabilities for each format (e.g., biological sequences).
- The benchmarks and evaluations of capabilities of the model, including considering the number of tasks it can perform without additional training, adaptability to learn new, distinct tasks, its level of autonomy and scalability, the tools to which it has access.
- Whether it has a high impact on the EU’s internal market due to its reach, which shall be presumed when it has been made available to at least 10,000 registered business users established in the EU.
- The number of registered end-users.

This creates a significant problem for any provider of a GPAI model, because even if the GPAI model does not appear to meet the criteria for having high-impact capabilities under Art.51(1)(a) and (2), it is possible that the Commission may, at any moment, simply decide ex officio to designate that GPAI model as having such capabilities notwithstanding. While Art.51(1)(b) requires the Commission to consider Annex XIII, the language used suggests there is room for the Commission to make ex officio designations of GPAI models with systemic risk without necessarily relying on the criteria set out in Annex XIII. Such designations trigger potentially significant compliance obligations for providers (see, in particular, Chapter 14).

Further, the Commission has the power to amend the criteria in Annex XIII, by means of a delegated act (e.g., if other features of GPAI models arise over time as these types of models advance generally).

Notification procedure – Providers can rely on the specific characteristics of their GPAI models to argue that there are no systemic risks.

Providers must continually assess the requirements of their GPAI models and keep up-to-date documentation as per the requirements of Annex XI (Art.53(1)(a) – see Chapter 13 for more information).

Where a provider determines that a GPAI model has high-impact capabilities, it must notify the Commission within two weeks that this threshold has been met and include the information necessary to demonstrate this (Art.52(1)). However, providers can, at the same time as making the notification, advance a case that, notwithstanding the fact the GPAI model meets the Art.51(1)(a) threshold, it nonetheless should not be classified as a GPAI model with systemic risk (i.e., even if the GPAI model ostensibly appears to have high-impact capabilities, providers can argue that specific characteristics or qualities of a particular GPAI model mean that systemic risk does not arise in the context of that model (Art. 52(2))).

However, the Commission may reject such arguments presented by providers, if it concludes that they are not “*sufficiently substantiated*” to prove that the specific characteristics of the GPAI model does not present systemic risk (Art.52(3)).

If the Commission becomes aware of a GPAI model that ostensibly presents systemic risks that has not been notified to the Commission, it may also decide to designate it as a GPAI model with systemic risk, in which case no express provision is made for providers to put forward their case against such a determination (Art.52(1)) (therefore likely encouraging over-inclusive notifications by providers).

Appeal procedure – Providers have unlimited opportunities to challenge designations made by the Commission.

If the Commission decides that a GPAI model has systemic risk (see Arts.51(1)(b) and 52(4) of the EU AI Act), the provider of that GPAI model can request a reassessment (at the earliest, six months after the Commission’s decision). The Commission must take the request into account and may re-assess the GPAI model against the criteria listed at Annex XIII (discussed above). Even after such re-assessment, providers have the ability to request another reassessment after a further six months – suggesting that providers have an unlimited number of reassessment requests at six-month intervals. However, any re-assessment request submitted by a provider should include compelling and novel arguments, incorporating details that have arisen in respect of the GPAI model since the date of the last decision by the Commission (indicating that a provider needs to present cogent arguments before its re-assessment request will be considered).

Context and illustrations



Commentary: High-impact capabilities

Providers must understand the term “*high-impact capabilities*” (Art.51(1)(a)) to be able to determine if a notification to the Commission is necessary. A literal interpretation of the term “*high-impact capabilities*” (see Rec. 111) suggests that a GPAI model must match or exceed the capabilities of the most advanced GPAI models. However, this interpretation leaves some ambiguity: The EU AI Act specifies that providers should use “*appropriate technical tools and methodologies, including indicators and benchmarks*” to evaluate these capabilities (Art.51(a)). Yet, it does not explain the criteria these tools should consider. Given how quickly new iterations of GPAI models are being deployed into the market, businesses will need to factor in the high pace of change in their assessments – a GPAI model that has a high impact right now (relative to the rest of the market) might not reach that same threshold a month later, as competitor GPAI models adapt and improve.

Additionally, there remains some ambiguity regarding the quantitative threshold specified in Art.51(2) for providers of GPAI models. While Art.51(2) states that GPAI models trained on computational resources exceeding 10^{25} FLOPs will be deemed to have “*high-impact capabilities*”, it does not set out how providers should measure FLOPs (e.g., do FLOPs used in the creation of prior iterations of a GPAI model count towards the total of the current iteration? If a GPAI model is branched, are FLOPs used to create each branch measured cumulatively or independently?). Although there is some guidance in Annex XIII of the EU AI Act on how the Commission may calculate a GPAI model’s training compute (e.g., by reference to variables such as the estimated cost of training or the estimated time required for training), it is ultimately unclear whether providers of GPAI models are required to use the same criteria. Even if they are, the use of these criteria does not answer the questions noted above.



Example: The notification procedure

Company A is a technology firm based in the United States and has developed a GPAI model that is trained with large volumes of financial data. Company B, a financial services provider based in Spain, builds an AI tool to forecast market trends and make investment decisions by integrating into Company A's GPAI model. This is done via an application programming interface so there is no other 1-to-1 interaction between Company A and Company B.

Company A (despite not being based in the EU) is subject to the EU AI Act (see Chapter 2) and will therefore need to carry out an assessment of whether its GPAI model has high-impact capabilities, with reference to Arts.51(1) (a) and (2). In the event that it does, Company A must notify the Commission within two weeks after reaching that conclusion, detailing its reasons for this determination (Art.52(1)).

In the meantime, Company B will need to keep track of any determinations or designations of systemic risk in relation to Company A's underlying GPAI model to assess how this might impact Company B's use of that GPAI model. If Company A did not notify the Commission, but was subsequently deemed by the Commission to be a provider of a GPAI model with systemic risk, this could also impact the operations and obligations of Company B.

As a result, the designation requirements in Arts.51 and 52 have the potential to significantly impact both companies that provide GPAI models (i.e., providers) and companies that use GPAI models provided by third parties (i.e., deployers) – see Chapter 11 for more information.



Commentary: The appeal procedure

Under the EU AI Act, providers of GPAI models have a significant procedural benefit when it comes to challenging the Commission's designations of GPAI models with significant risk. If the Commission decides that a GPAI model presents systemic risk under Art.51(1)(b), the provider of that GPAI model can request a reassessment of this decision after six months. The Commission is then obligated to reassess the GPAI model based on the criteria listed in Annex XIII.

This reassessment process is not a one-time opportunity. Providers can request another reassessment six months after receiving the Commission's decision, and this cycle can potentially continue indefinitely (at least for as long as the provider is able to present cogent arguments in its re-assessment request). This essentially grants providers an unlimited number of opportunities (at six-month intervals) to challenge the Commission's designation of systemic risk for as long as the provider can make arguments regarding that designation.

This suggests that providers may have a greater degree of influence over the status of a GPAI model than initially appears, notwithstanding the Commission's ability to make designation decisions ex officio or on referral. It further highlights how a deep technical understanding of GPAI models (and their iterations) will benefit providers most when it comes to challenging designations of systemic risk.

Chapter 13

GPAI models – General obligations of providers of GPAI models

Executive summary

Under the EU AI Act, providers of GPAI models must prepare and maintain extensive technical information and documentation relating to the training and functionality of such GPAI models.

Providers of GPAI models will be expected to disclose different degrees of information to: (i) the AI Office and national competent authorities (upon request); (ii) other downstream providers of AI systems; and (iii) the wider public.

Providers of open-source GPAI models can benefit from a limited exemption to some of these compliance and disclosure obligations.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.53 – 54** of the EU AI Act which set out the obligations for providers of GPAI models. These include: (i) creating and maintaining technical documentation relating to their GPAI models; (ii) sharing information with downstream providers of AI systems; (iii) complying with EU copyright law; and (iv) (for providers in third countries) appointing an authorised representative in the EEA to verify the provider's compliance with its obligations under (i) – (iii).

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model or that has an AI system or a GPAI model developed, and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
 - **Downstream provider** (Art.3(68)) – A provider of an AI system, including a GPAI system, which integrates an AI model, regardless of whether the AI model is provided by the provider itself and vertically integrated or provided by another entity based on contractual relations.
 - **GPAI system** (Art.3(66)) – An AI system which is based on a GPAI model and which has the capability to serve a variety of purposes, both for direct use as well as for integration in other AI systems.
 - **GPAI model** (Recs.97-99, Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market, and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.
- A full list of defined terms can be found in the [Glossary](#).

Analysis

Technical compliance documentation

Providers must provide and maintain specific documentation relating to their GPAI models (Art.53(1)(a)). These may be requested by the AI Office and/or national competent authorities in the relevant EEA state in which the provider has made the GPAI model accessible. At a minimum, the documentation should include the following information (Annex XI):

- A general description of the GPAI model which includes detailed reference to the following characteristics:
 - The tasks that the GPAI model can perform.
 - The type and nature of AI systems that the GPAI model can be applied to.
 - Applicable acceptable use policies.
 - The date of release.
 - The methods of distribution.

- The number of parameters and general architecture of the GPAI model.
 - The modality and type of input and output (e.g., text, image, sound, etc.).
 - The licence that governs the use of the GPAI model (as noted below, Art.53(2) exempts GPAI models provided under free or open-source licences from certain obligations).
- Information on the process of development of the GPAI model, including but without limitation:
- The technical requirements for the GPAI model to be integrated into an AI system (e.g., instructions of use, infrastructure, or tools).
 - The methodologies, techniques, assumptions, optimisations, design specifications, and choices used for training the GPAI model.
 - Information on the training, testing, and validation data which includes the type of data, where the data came from, how the data was curated (e.g., cleaning, filtering, etc.), the number of data points and what they relate to, any methods applied to detect relevant biases, and any measures implemented to assess the unsuitability of data sources.
 - The computational resources used for training purposes, the training time, and the known or estimated energy consumption of the GPAI model.
- How the GPAI model currently or potentially interacts with external hardware or software.
- The versions of relevant software related to the use of the GPAI model.
- Information on the elements of the GPAI model and the process of development, including but without limitation:
- The technical requirements for the GPAI model to be integrated into an AI system (e.g., instructions of use, infrastructure, or tools).
 - The modality (e.g., text, image, etc.) and format of the inputs and outputs and their maximum size (e.g., context window length, etc.).
 - Information on the training, testing, and validation data which includes the type of data, where the data came from, and how the data was curated (e.g., cleaning, filtering, etc.), where applicable.

Compliance with the requirements above is subject to any relevant intellectual property rights, confidential business information, or trade secrets that may affect such providers of GPAI models (Art.53(1)(b)).

Copyright compliance

Providers of GPAI models must implement a policy to comply with EU law on copyright and related rights (Art.53(1)(c)). Specifically, this is to identify and comply with rightsholders' reservations of rights (or "opt-outs") in relation to the text and data mining exception to copyright and database rights under the Copyright in the Digital Single Market Directive. Providers may use state-of-the-art technologies to help achieve compliance (Art.53(1)(c)). Pursuant to the GPAI Code of Practice (currently in its third draft at the time of writing), signatories to that Code commit to reproduce and extract only lawfully accessible copyright-protected content, and not to circumvent copyright protection measures, and make a reasonable effort to exclude websites that publish copyright-infringing materials from their training data (though the Code does not require a work-by-work review to avoid accidental infringement of copyright). Signatories to the Code are required to prohibit copyright-infringing uses of their GPAI models in their acceptable use policy. Finally, signatories are required to establish a clear point of contact and complaint mechanism for rightsholders wishing to raise concerns about potential non-compliance.

Information to assist downstream providers

Providers must prepare, maintain, and make available information and documentation relating to their GPAI models to other downstream providers. This information must be sufficient to provide them with a good understanding of the capabilities and limitations of the GPAI model. This enables downstream providers to integrate such GPAI models into their own AI systems and comply with their own obligations under the EU AI Act (Art.53(1)(b)(ii)). In addition to providing a general description of the GPAI model as required by Art.53(1)(a) set out above, providers must provide the following information (Art.53(1)(b)(iii) and Annex XII):

Transparency information

Providers must publish transparency information, including a detailed summary of all the material used to train the GPAI model. This summary must be prepared on the template document provided by the AI Office (Art.53(1)(d)). Additional transparency obligations apply in relation to AI systems and are discussed in Chapters 7 (in relation to high-risk AI systems) and 11 (for certain AI systems).

Authorised representatives

Providers of GPAI models established in third countries must appoint an authorised representative which is established in the EEA before placing a GPAI model on the EEA market. The authorised representative will be responsible for verifying the provider's compliance with its obligations under Art.53, maintaining compliance documentation for up to ten years from the point of placing the GPAI model on the EEA market, and cooperating with the AI Office and national competent authorities as required.

Authorised representatives may liaise with the AI Office and national competent authorities in addition to, or on behalf of, providers of GPAI models (depending on the terms of their written appointment). They must terminate the appointment if they have reason to believe that the relevant GPAI model provider is acting contrary to its obligations under Art.53 and inform the AI Office immediately.

Exemptions

Providers of open-source GPAI models are exempt from some of these requirements, namely to: (i) prepare technical documentation for the AI Office or national competent authorities; (ii) provide integration instructions to downstream providers; or (iii) appoint an authorised representative (Arts.53(2) and 54(6)).

A GPAI model is considered open source if it is released under a free and open-source licence that permits access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available (Art.54(6)).

However, providers of open-source GPAI models must still comply with applicable copyright law and must publish a summary on the training materials used for the public.

In accordance with the GPAI Code of Practice, signatories who are providers of free open-source GPAI models do not have to prohibit copyright-infringing uses in their acceptable use policy.

Context and illustrations



Commentary: Technical compliance documentation

Providers must maintain technical compliance documentation for their GPAI models, which may be requested by the AI Office or national competent authorities within the relevant EEA state where the GPAI model is accessible.

A literal interpretation of this requirement suggests that records must be kept on various aspects of the GPAI model (Annex XI). This includes a general description of the model, such as the tasks it can perform, the types of AI systems it can be applied to, applicable acceptable use policies, the date of release, methods of distribution, the number of parameters, and its general architecture. However, this interpretation leaves some ambiguity, particularly regarding the exact level of detail needed for compliance.

To address this, the EU AI Act specifies that providers should include information on the modality and type of input and output (e.g., text, image, sound), the type of licence, and the process of development. Providers must also detail the methodologies, techniques, assumptions, optimisations, design specifications, and choices used during the training of the GPAI model.

By keeping comprehensive and up-to-date documentation, businesses can make a strong case for compliance with the EU AI Act and be prepared for any inquiries from national competent authorities.



Example: Downstream providers

Providers must prepare, maintain, and share information about their GPAI models with downstream providers. This helps downstream providers understand the model's capabilities and limitations, allowing for proper integration into their AI systems and compliance with the EU AI Act (Art.53(1)(b)(ii)).

For example, Company X in Germany has developed a GPAI model for voice-to-text natural language processing. Company Y, a customer service provider in Italy, uses this model to improve customer interactions. Company X must give Company Y detailed documentation, including a general description of the GPAI model, the tasks it can perform, and the types of AI systems it can be applied to (Art.53(1)(a)).

Company X should also be mindful of intellectual property rights and trade secrets when sharing this information (Art.53(1)(b)). By providing comprehensive and accurate documentation, Company X ensures that Company Y can integrate the GPAI model effectively and remain compliant with the EU AI Act.



Example: Exemptions

Providers of open-source GPAI models are exempt from some requirements under the EU AI Act. Specifically, they are not required to prepare technical documentation for the AI Office or national competent authorities, provide integration instructions to downstream providers, or appoint an EEA representative (Arts.53(2) and 54(6)).

For example, if Company Z develops an open-source GPAI model for image recognition, it does not need to create detailed technical documentation for regulatory bodies or provide integration instructions to other companies using its model. It also does not need to appoint a representative within the EEA. But Company Z must ensure that the model is released under an open-source licence and that all relevant model parameters and architecture information are publicly accessible. Additionally, it must provide a summary of the data used for training the model, ensuring transparency and compliance with copyright regulations.

By understanding these exemptions, businesses can effectively navigate their obligations and leverage the benefits of open-source development while remaining compliant with the EU AI Act.

Chapter 14

GPAI models – General obligations of providers of GPAI models with systemic risk

Executive summary

All providers of GPAI models are subject to the general compliance obligations set out in Chapter 13. Providers of GPAI models with “systemic risk” are subject to additional obligations to identify, assess, prevent, and/or mitigate existing and potential systemic risks in their GPAI model. There is no one-size-fits-all approach to compliance with these obligations. Providers of GPAI models with systemic risk will need to consider how best to ensure compliance, taking into account the “state of the art”.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.55 and Annex XI** of the EU AI Act which set out the obligations for providers of GPAI models with systemic risk. These obligations include: (i) performing GPAI model evaluations; (ii) assessing and mitigating potential systemic risk; (iii) documenting and reporting incidents; (iv) implementing adequate cybersecurity protection for their GPAI model; and (v) providing more extensive technical documentation.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model or that has an AI system or a GPAI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
- **GPAI model** (Recs.97 – 99; Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of

competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.

- **GPAI model with systemic risk** (Art.51) – A GPAI model that meets the conditions of Art.51.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Overview

As set out in Chapter 12, some GPAI models are deemed to have “systemic risk”. All providers of GPAI models are subject to the general compliance obligations set out in Chapter 13 (including obligations to provide technical information and documentation, and obligations to disclose certain information to regulators and the public). Providers of GPAI models with systemic risk are subject to additional obligations, as set out below.

Technical compliance documentation

In addition to the compliance documentation required for all GPAI models under Arts.53 and 54, providers of GPAI models with systemic risk must prepare and maintain the following information (Section 2 of Annex XI):

- A detailed description of the evaluation strategies, including evaluation criteria, metrics, and methodology on the identification of limitations as well as overall evaluation results.
- Where applicable, a detailed description of the measures put in place to conduct internal and/or external adversarial testing (e.g., red teaming), model adaptations, alignment, and fine-tuning.
- Where applicable, a detailed description of the system architecture explaining how software components build or feed into each other and integrate into the overall processing.

Providers of GPAI models with systemic risk must provide this information upon request to the AI Office and/or national competent authorities (Art.54(3)(b)).

GPAI model evaluations

Providers of GPAI models with systemic risk must use standard protocols and tools that reflect the state-of-the-art practices to perform and document evaluations of the GPAI models with systemic risk (Art.55(1)(a)). Examples include carrying out and documenting adversarial testing – i.e., a method of testing the robustness and security of a GPAI model by providing atypical inputs, for the purpose of identifying and resolving risks that could arise from unexpected (or “*adversarial*”) use of the GPAI model. Providers of GPAI models with systemic risk should carry out GPAI model evaluations in accordance with relevant standardised protocols and state-of-the-art tools (to the extent available). The core purpose of each evaluation should be to test the relevant GPAI model’s risk levels as well as identify and mitigate the systemic risks posed by that GPAI model.

Risk assessments

Providers of GPAI models with systemic risk must assess and mitigate existing and potential risks that may arise from the development, commercialisation, or use of the GPAI model (Art.55(1)(b)). This imposes a continuing obligation on providers not only to address existing risks but also proactively identify potential risks that may likely arise in the future. In practical terms, this may require implementing risk management policies, such as accountability and governance processes, and taking measures throughout the model’s entire lifecycle, starting with development. It also requires monitoring the GPAI model after placing it on the market and cooperating with actors along the AI value chain (e.g., downstream providers of AI systems (Rec.114)).

Incident reporting obligations

Providers of GPAI models with systemic risk must also keep track of, document, and report serious incidents and possible corrective measures to the AI Office and, where appropriate, national competent authorities without delay (Art.55(1)(c)). A “serious incident” includes (Rec.155 and Art.3(49)):

- An incident or malfunctioning leading to death or serious damage to health.
- A serious and irreversible disruption of the management and operation of critical infrastructure.
- Infringements of fundamental rights.
- A serious damage to the property or the environment.

Cybersecurity obligations

Providers of GPAI models with systemic risk established in third countries must implement an adequate level of cybersecurity protection that protects both the GPAI model and the physical infrastructure on which the GPAI model runs (Art.55(1)(d)).

Cybersecurity protection for GPAI models with systemic risk should safeguard against the following risks (Rec.115): accidental model leakage, unauthorised releases, circumvention of security measures, cyberattacks, unauthorised access, and “*model theft*”.

Appropriate cybersecurity measures against the systemic risks associated with malicious use or malicious attacks include: (i) implementing technical and organisational security measures to protect model weights, algorithms, servers, and data sets; (ii) establishing appropriate cybersecurity policies and adequate technical solutions; and (iii) setting up cyber and physical access controls, in each case as appropriate to the relevant circumstances and the risks involved (Rec.115).

Demonstrating compliance

Art.55(2) allows providers of GPAI models with systemic risk to demonstrate compliance with their cybersecurity obligations under Art.55(1)(d) (explained above) by following codes of practice (Art.56 – see Chapter 15) and/or EU harmonised standards (Art.40 – see Chapter 13) where available. Conforming to codes of practice/ harmonised standards creates a presumption of conformity with Art.55(1) – provided that such codes of practice/ harmonised standards cover the requirements of Art.55(1).

Where the provider of a GPAI model with systemic risk chooses not to adhere to an approved code of practice/ harmonised standard, that provider does not benefit from the presumption of conformity noted above. The provider must demonstrate alternative adequate means of compliance, and the Commission has the power to assess those means of compliance (Art.55(2)).

Confidentiality

Under Art.78, the Commission and the national competent authorities are required to respect the confidentiality of information and documentation disclosed by providers of GPAI models with systemic risk under Art.55 (including trade secrets or intellectual property) subject to the exceptions noted in Art.78 (Art.55(3)).

Context and illustrations



Commentary: Overlapping cybersecurity regimes

Providers of GPAI models with systemic risk must implement cybersecurity protection to safeguard the model and its physical infrastructure, in particular from malicious use or malicious attacks (Art.55(1)(d)). Rec.115 provides further detail and helpful guidance to providers on the recommended types of protective measures that they should look to implement. However, some providers may also be caught by cybersecurity obligations imposed by the NIS 2 Directive on cybersecurity. Where a business falls within the scope of NIS 2 and is also a provider of a GPAI model with systemic risk, that business may have to report the same incident under both Art.55(1)(c) of the EU AI Act and Art.6(6) NIS 2/Art.23(3) NIS 2, as well as the Implementing Regulation to the NIS 2 Directive. In addition, where an incident affecting a GPAI model with systemic risk also involves the processing of personal data, it is possible that that incident would also be reportable under Arts.33 and 34 GDPR.

As a result, businesses that are providers of GPAI models with systemic risk may be faced with parallel reporting obligations under multiple EU laws, in relation to the same incident. There is no easy way to avoid this challenge, and businesses that provide GPAI models with systemic risk will need to keep a close eye on their incident reporting obligations.



Example: Risk assessment during a GPAI model's entire lifecycle

For example, a provider of a GPAI model with systemic risk will need to start the process of risk assessments while the GPAI model is still in development. If the GPAI model has multiple iterations where its functionalities significantly change from the previous version, the provider will need to carry out further risk assessments to factor in each such change.

Prior to placing the GPAI model with systemic risk on the market for adoption by end-users and downstream providers, the provider will also need to carry out a further risk assessment. Once the GPAI model with systemic risk is on the market, the provider will need to continue to monitor the GPAI model for potential or new risks.

In practice, this will likely require ongoing communications with actors along the AI value chain (e.g., at least some downstream providers that adopt the GPAI model with systemic risk, and potentially even end-users). In other words, the provider will need to continue carrying out risk assessments during the GPAI model's entire lifecycle.



Commentary: "Model theft" and "state of the art"

In several places, the terminology used in the AI Act is unclear, for example:

"Model theft": Providers of GPAI models with systemic risk are required to protect against "*model theft*", a term that is only used once in the EU AI Act (Rec.115) and is not defined. Academic texts suggest that "*model theft*" includes the scenario where a third party with legitimate access to the GPAI model but "*no prior knowledge of [its...] parameters or training data, aims to duplicate the functionality of (i.e., 'steal') the model.*" Therefore, providers of GPAI models with systemic risk must protect against classic cybersecurity challenges, but also against reverse-engineering based on otherwise lawful access to the GPAI model.

"State of the art": Art.55(1)(a) requires a GPAI model evaluation to be carried out using protocols and tools that reflect the "*state of the art*" – a concept always in flux. Providers are effectively responsible for determining what is "*state of the art*" in their particular context (given that the state of the art during the development of a GPAI model is likely to change over the lifespan of most GPAI models). "*State of the art*" is used in several places in the EU AI Act and introduces ambiguity that makes compliance challenging, as the target is always moving. It also risks regulators looking at measures with the benefit of hindsight to determine what was "*state of the art*" at the time. It is possible that with time, what constitutes "*state of the art*" will develop. But until then, providers of GPAI models with systemic risk should carefully consider what "*state of the art*" means and be prepared to defend the choices they make.

Chapter 15

GPAI models – Codes of practice

Executive summary

The EU AI Act provides for the creation of codes of practice, which can be used by GPAI model providers to demonstrate compliance with the EU AI Act.

Providers may rely on codes of practice to show compliance with Art.53 (documentation) and Art.55 (systemic risk obligations) until harmonised standards are adopted.

The drafting of the first code of practice has started, involving GPAI model providers, civil society organisations, industry, academia, and other relevant stakeholders.

The final version of the first code of practice is expected to be published by August 2025. If adequate codes of practice covering Arts.53 and 55 are not agreed by August 2025, the Commission can instead adopt implementing acts setting out common rules for compliance.

Relevant sections of the EU AI Act

This Chapter focuses on **Art.56** of the EU AI Act – specifically, the role, creation process, and content of the codes of practice. It also provides insights on **Arts.53** and **55** of the EU AI Act as there is a presumption of conformity for GPAI model providers who comply with the codes of practice until harmonised standards are adopted.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **GPAI model** (Recs.97 – 99; Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market, and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.
- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model or that has an AI system or a GPAI model developed, and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
- **AI Office** (Arts.3(47) and 64) – The Commission's function of contributing to the implementation, monitoring, and supervision of AI systems and GPAI models, as well as AI governance. The EU AI Act states that references to the AI Office should therefore be construed as references to the Commission.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Need for codes of practice – The EU AI Act provides for the adoption of codes of practice, intended to bridge the gap between GPAI model provider obligations that enter into effect from August 2025, and the later adoption of harmonised standards.

The EU AI Act places thorough obligations on GPAI model providers (Arts.53, 54, and 55), coming into effect from August 2025 (Art.113). Pursuant to the compliance model of the EU AI Act, GPAI model providers must demonstrate their conformity with such obligations.

The EU AI Act provides for a specific presumption of conformity: Compliance with European harmonised standards grants providers a presumption of conformity to the extent that such standards cover those obligations (Art.53).

However, because the standardisation process can be lengthy, the EU AI Act also provides for an interim solution. Until a harmonised standard is published, compliance with codes of practice will grant GPAI model providers the presumption of conformity. Failing which, they will need to demonstrate alternative adequate means of compliance for assessment by the Commission (Art.53(4)).

Encouragement of codes of practice – The AI Office encourages and facilitates the creation of relevant codes of practice.

In a press release dated 30 July 2024, the AI Office confirmed that it would facilitate an iterative drafting process to ensure that the codes of practice effectively address the provisions under the EU AI Act, duly reflect the state of the art, and take into account a diverse set of perspectives (see also Rec.116).

To that end, in September 2024 the AI Office launched a multi-stakeholder consultation regarding what the first code of practice should cover and held an opening plenary session with nearly 1,000 participants. Eligible participants included:

- **GPAI model providers** with existing or planned operations in the EU.
- **Downstream providers and other industry organisations** with existing operations or planned operations in the EU, and a legitimate interest.
- **Academia, independent experts, and related organisations** with relevant expertise in at least one topic covered by the working groups.
- Other stakeholder organisations (including civil society organisations representing a specific stakeholder group or organisations representing rightsholders) with a presence in the EU, legitimate interest, and able to demonstrate that they are representative for the relevant stakeholder group.

Going forward, four working groups have been established, reflecting the categories of provisions under the GPAI model section of the EU AI Act, namely:

- **Working Group 1 – Transparency and copyright-related rules** – This working group aims at detailing out documentation to downstream providers and the AI Office on the basis of Annexes XI and XII to the EU

AI Act, preparing policies to be put in place to comply with EU law on copyright and related rights, and making publicly available a summary about the training content.

- **Working Group 2 – Risk identification and assessment measures for systemic risks** – This working group aims at detailing the risk taxonomy based on a proposal by the AI Office, and identifying and detailing relevant technical risk assessment measures, including model evaluation and adversarial testing.
- **Working Group 3 – Risk mitigation measures for systemic risks** – This working group aims at identifying and detailing relevant technical risk mitigation measures, including cybersecurity protection for the GPAI model and the physical infrastructure of the model.
- **Working Group 4 – Internal risk management and governance for GPAI model providers** – This working group aims at identifying and detailing policies and procedures to operationalise risk management in internal governance of GPAI model providers, including keeping track of, documenting, and reporting serious incidents and possible corrective measures.

These working groups are chaired by independent experts appointed by the AI Office and are responsible for synthesising submissions from the plenary participants. The list of such chairs and vice-chairs was published by the Commission on 30 September 2024 during the kick-off plenary meeting and predominantly features profiles from academia and research institutions.

The drafting process has been iterative through plenary sessions and various workshops to which GPAI model providers were invited for their input.

The AI Office will ensure transparency into these discussions, such as by drawing-up meeting minutes and making these available to all plenary participants.

The final version of the first code of practice is expected by August 2025, to be presented in a closing plenary session (Art.56(9)).

Approval of codes of practice – Further to the assessment by the AI Office and the AI Board, the Commission may approve the draft codes of practice.

Once the final draft is issued, the AI Office and the AI Board will assess the adequacy of the code of practice at issue and will publish their assessment.

Following that assessment, the Commission may decide to approve such code and give it general validity within the EEA by means of implementing acts (Art.56(6)).

If by 2 August 2025 (i.e., the time the EU AI Act becomes applicable) the first code of practice is not finalised or deemed adequate by the AI Office, the Commission may provide common rules for the implementation of the relevant obligations (Art.56(9)).

Content of codes of practice – The codes of practice should cover obligations provided for GPAI model providers in the EU AI Act, but their specific content is yet to be defined.

Codes of practice should cover obligations for GPAI model providers, including where they present systemic risks. In addition, as regards systemic risks, codes of practice should help to establish a risk taxonomy of the type and nature of the systemic risks at EEA level, including their sources. Codes of practice should also be focused on specific risk assessment and mitigation measures (Art.56(2); Rec.116).

In codes of practice in other sectors, relevant signatories have agreed to high-level commitments, qualitative reporting elements, service level indicators, key performance indicators, or other structural indicators specifically defined under such codes.

It seems that the AI Office took the same approach for the first code of practice. Three drafts have already been made available to the public ([first draft](#), [second draft](#), and [third draft](#)), addressing key considerations for GPAI model providers when complying with Chapter V of the EU AI Act. In particular, the third draft of the code published on 11 March 2025 is split out into four different documents to correspond with the following areas of compliance for GPAI model providers: (i) a general overview of the commitments; (ii) transparency; (iii) copyright; and (iv) safety and security.

Overall commitments: This section of the code provides an overview of all the relevant commitments related to the transparency, copyright, and safety and security compliance obligations under Art.53. Notably, it separates the commitments by general GPAI model providers (in relation to transparency and copyright obligations) from the commitments by general GPAI model providers with systemic risk (in relation to safety and security obligations).

Transparency: This section of the code includes a model documentation form to help GPAI model providers meet their compliance obligations in relation to keeping up-to-date model information for the benefit of downstream providers and/or the AI Office under Art.53(1). The relevant sections of the form include general information about the GPAI model; model properties; methods of distribution and licences; training process; information on the data used for training, testing, and validation; computational resources; and additional information required from providers of the GPAI model with systemic risk.

Copyright: This section of the code establishes appropriate measures for GPAI model providers to follow in order to demonstrate compliance with the commitment to implement and keep up-to-date a copyright policy in accordance with Art.53(1)(c).

Safety and security: This section of the code only applies to GPAI model providers with systemic risk (as opposed to all GPAI model providers). It details the relevant commitments and measures that they must meet to establish compliance with their obligations related to the continuous assessment and mitigation of systemic risks.

Review and adaptation of codes of practice – The AI Office will encourage and facilitate the review and adaptation of the codes of practice, particularly in light of emerging standards.

Because these codes of practice are part of a quickly changing technological and legal landscape, they cannot remain still and must be able to evolve with technological, societal, market, and legislative developments. To that end, the EU AI Act requires the AI Office to encourage and facilitate the review and adaptation of the codes of practice (Rec.116; Art.56(8)).

Currently, nothing has been announced regarding processes for post-codes of practice monitoring and updating. Following the first iteration, certain stakeholders contemplate the setting-up of a permanent taskforce, featuring signatories of the codes and other oversight bodies, to review and adapt such codes based on technological, societal, market, and legislative developments.

Context and illustrations



Commentary: Presumption of conformity under Art.53(4)

Given the breadth of documentation required under Art.53, the presumption of conformity under Art.53(4) should help providers by providing them with a clear route to compliance. It is clear from the first sentence of Art.53(4) that compliance with the code of practice will trigger the presumption of conformity prior to the publication of harmonised standards.

However, once a harmonised standard is published, it is slightly less clear whether continued compliance with a code of practice (as opposed to the published harmonised standard) would continue to trigger the presumption of compliance. In practice, we would expect the codes of practice to be updated to ensure that they align with newly published standards, in line with Art.56(8).



Commentary: One or several codes of practice?

The wording of Art.56 is clear in that it provides for the drawing-up of “*codes of practice*”, plural. The EU AI Act specifies that “*the codes of practice should represent a central tool for the proper compliance with the obligations provided for under the [EU AI Act for GPAI model providers]*” (Rec.117).

Nonetheless, currently, the AI Office’s approach has been focused on preparation of the “*first General-Purpose AI Code of Practice*”. There is no clear indication at this stage on whether this first code of practice will be followed by others, and, if so, whether such codes will concern specific obligations under the EU AI Act, specific activity sectors, or will rather be successive versions of the same code as updated based on technological, societal, market, and legislative developments.



Commentary: Legal significance of codes of practice under Art.56(7)

While Art.56(3) clearly positions adherence to codes of practice as voluntary, sophisticated GPAI model providers should consider that Art.56(7) introduces a nuanced regulatory dynamic by allowing the AI Office to directly “invite” GPAI model providers to adhere to these codes. Sophisticated GPAI model providers should recognise that regulatory invitations under Art.56(7) may create implicit expectations of compliance, especially for GPAI model providers whose adherence might be limited to Art.53 obligations (e.g., if their models do not present systemic risks).

Regulatory authorities or national courts could reference these codes as indicative evidence of prevailing industry standards, potentially influencing their assessments of compliance. Providers of GPAI models should therefore carefully evaluate their adherence decisions, understanding that voluntary codes of practice may, in fact, carry considerable practical and legal significance.

Chapter 16

Measures designed to support innovation

Executive summary

The EU AI Act contains various measures to facilitate the development and testing of AI systems, through regulatory sandboxes and real-world testing, before going to market. The aim of these measures is to encourage innovation, although it remains unclear whether that outcome will be achieved.

Providers participating in a sandbox are not subject to administrative fines under the EU AI Act (provided they adhere to the rules governing the sandbox). However, those providers may still be liable to claims for damages from third parties for harm arising from the use of the sandbox.

Real-world testing conducted outside sandboxes lets providers take a greater degree of control but provides no direct protection against liability or administrative fines.

Relevant sections of the EU AI Act

This Chapter focuses on Arts.57 – 63 of the EU AI Act – specifically, the measures to facilitate the development and testing of AI systems within (and outside of) national AI regulatory sandboxes.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **AI regulatory sandbox** (Art.3(55)) – A controlled framework set up by a competent Member State authority which offers providers of AI systems the possibility to develop, train, validate, and test, where appropriate in real-world conditions, an innovative AI system, pursuant to a sandbox plan for a limited time under regulatory supervision.
- **Testing in real-world conditions** (Art.3(57)) – The temporary testing of an AI system for its intended purpose in real-world conditions outside a laboratory or otherwise simulated environment, with a view to gathering reliable and robust data and to assessing and verifying the conformity of the AI system with the requirements of the EU AI Act. It does not qualify as placing the AI system on the market or putting it into service within the meaning of the EU AI Act, provided that all the conditions laid down in Arts.57 or 60 are fulfilled.
- **Real-world testing plan** (Art.3(53)) – A document that describes the objectives; methodology; geographical, population, and temporal scope; monitoring; organisation; and conduct of testing in real-world conditions.
- **Sandbox plan** (Art.3(54)) – A document agreed between the participating provider and the competent authority describing the objectives, conditions, timeframe, methodology, and requirements for the activities carried out within a sandbox.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Establishment of national AI regulatory sandboxes

(Recs.138 and 139; Arts.57 – 59) – The EU AI Act introduces a framework for developing and testing AI systems in a controlled environment called an “AI regulatory sandbox”. Member States are required to establish at least one national AI regulatory sandbox for providers to participate in. A sandbox must be created in line with detailed implementing acts to be adopted by the Commission in the future.

- **Sandboxes** – The concept of sandboxes exists in many other areas of EU law (including data protection, blockchain, and so on). The idea is to provide a controlled environment, set up by regulators, that allows organisations to test innovative technologies, products, or services (like AI systems) under more relaxed regulatory conditions, while ensuring appropriate safeguards and oversight. Because sandboxes necessarily cover the development phase of AI systems, in this Chapter, the term “provider” is used to refer to both: (i) businesses that have already placed their AI systems onto the market or put them into service; and (ii) businesses that are still developing their AI systems.
- **Compulsory sandboxes** (Art.57(1) and (4)) – Each Member State is required to: (i) establish alone or in conjunction with other Member States, at least one sandbox at national level that is operational by 2 August 2026; (ii) allocate sufficient resources, including guidance, supervision, and support, for effective and timely compliance; and (iii) ensure cooperation between its national competent authorities and other authorities within the sandbox ecosystem.
- **Optional sandboxes** (Art.57(2) and (3)) – Member States may establish additional sandboxes, either alone or jointly with other Member States. In addition, the EDPS may also establish a sandbox for EU institutions, bodies, offices, and agencies and will have roles and responsibilities akin to that of a national competent authority.
- **Objectives of sandboxes** (Art.57(9)) – Sandboxes should serve a limited list of purposes: legal certainty to improve regulatory compliance of AI systems and evidence-based regulatory learning; sharing best practices; fostering innovation and competitiveness and facilitating the development of an AI ecosystem; and accelerating access to the EU market for AI system.
- **Reporting obligations** (Art.57(7) and (16)) – Upon request of a provider, the competent authorities are required to provide a written proof of activities successfully carried out in the sandbox. In addition, competent authorities shall provide exit reports detailing the activities carried out in the sandbox and the related results and learning outcomes, as well as annual reports on the implementation of the sandboxes and outcomes. Access to the exit reports can be restricted but abstracts of the annual reports shall be public.
- **Surveillance of sandboxes** (Art.57(10-11)) – The competent authority is required to involve other relevant authorities (e.g., data protection authorities where the AI system involves processing of personal data) in the supervision of the sandbox. Their roles and corrective powers are not affected by the operation of AI systems in the sandbox.
- **Sandbox participants** – The rules governing access to sandboxes will be further set out in implementing acts to be adopted by the Commission (Art.58(1)). However, Art.58(2) outlines principles that will apply to sandboxes, including: openness to any provider fulfilling selection criteria; broad and equal access; and free access for SMEs.
- **Future guidance for sandboxes** (Art.58) – The Commission will adopt implementing acts and guidance, which will specify detailed arrangements for the establishment, development, implementation, operation, and supervision of the sandboxes.

Obligations on sandbox participants (Arts.3(54), 57(5), 57(12), and 63(1)) – Providers of AI systems who participate in a sandbox are required to adhere to agreed plans for participation, and remain liable for any damages to third parties that may arise as a result of their participation in a sandbox, subject to some exceptions as set out below:

- **Obligation on sandbox participants** (Arts.54(3) and 57(5)) – Sandbox participants are required to agree to a specific “*sandbox plan*” with the applicable national competent authority which describes the objectives, conditions, timeframe, methodology, and requirements for the activities to be carried out within that sandbox.
- **Liability of sandbox participants** (Art.57(12)) – Sandbox participants are liable under applicable EU and Member State law for any damage to third parties that arises out of the development and testing that they conduct within a sandbox. However, administrative fines under the EU AI Act will not be imposed if that sandbox participant complied with its sandbox plan and applicable terms and conditions, and participated in good faith according to the guidance given by the applicable national competent authority.

Processing personal data within AI regulatory sandboxes (Rec.140; Arts.57(10) and 59) – The EU AI Act permits the processing of personal data for the purpose developing, training, and testing public interest-focussed AI systems.

Where an AI system is being developed for a purpose that is in the public interest, and that AI system is being developed, trained, and/or tested within a sandbox, the EU AI Act provides a legal basis for the processing of personal data that was originally collected for a separate purpose, under Art.6(4) GDPR (Rec.140 EU AI Act). Rec.50 GDPR explains that where “*processing is necessary for the performance of a task carried out in the public interest [...EU] law may determine and specify the tasks and purposes for which the further processing should be regarded as compatible and lawful.*”

To simplify:

- Art.5(1)(b) GDPR sets a general rule that personal data cannot be processed for a new purpose that is incompatible with the original purpose for which the personal data was processed.
- Art.6(4) GDPR sets out the conditions and criteria under which further processing personal data for a new purpose can be lawful and compatible.
- Rec.140 EU AI Act clarifies that further processing of personal data for the purposes of developing, training, and testing AI system in a sandbox context, for public interest purposes, is permissible under the GDPR, provided that appropriate safeguards are followed.
- Art.59(1) EU AI Act clarifies the scope of this permission and explains the applicable limits.

Purposes that are deemed to be within the public interest in the context of sandboxes include public safety, health, energy sustainability, and transport safety, among others (Art.59(1)(a)).

The processing of personal data within a sandbox is still subject to the usual rules set out in the GDPR (i.e., providers will still need to give effect to the rights of data subjects, maintain records of processing, complete data protection impact assessments where required, etc.).

Testing of high-risk AI systems in real-world conditions outside of AI regulatory sandboxes (Arts.60 – 61)

– The EU AI Act allows providers of certain high-risk AI systems to test their AI system outside of a sandbox environment, in “*real-world conditions*”, before being placed on the market. Providers of eligible high-risk AI system must comply with the following (among others):

- **Eligible high-risk AI system** (Art.60(1) and Annex III) – The applicable high-risk AI system must be intended to be used for biometrics; critical infrastructure; education and vocational training; employment; accessing essential private and public services and benefits; law enforcement; and migration and administration of justice. Systems that are deemed high-risk under Art.6(1) are not eligible (e.g., AI systems that fulfil the “safety-critical AI systems” test – see Chapter 6).
- **Testing plan** (Art.60(1), (4)(a)-(b), (4)(f), and (4)(k)) – Providers must set up a testing plan that outlines, among other things, the objectives and conduct of testing in real-world conditions and have it approved by the responsible market surveillance authority. The duration of testing should last no longer than six months, and the provider must be able to reverse or disregard the predictions, recommendations, or decisions made by the AI system during that period. The Commission will adopt implementing acts to specify the required contents of the plan.
- **Informed consent** (Arts.60(4)(i), (5), and 61) – Where individuals are the subjects of testing in real-world conditions, providers must obtain documented, freely given informed consent from those individuals prior to their participation in the testing (subject to a narrow exception concerning law enforcement). This consent can be revoked, in which case the provider is required to immediately cease using that individual’s data for testing purposes.

- **Other requirements** (Art.60(4)(c)-(e)) – Providers must, among other things, register their high-risk AI system, be established in the EU or appoint a legal representative established in the EU, and restrict the transfer of data to non-EU countries. The provider’s registration details are generally only accessible to market surveillance authorities and the Commission (Art.71(4) – see Chapter 18).

A provider of a high-risk AI system tested in real-world conditions outside sandboxes will be liable for any damage caused during the testing phase. Unlike sandbox testing, complying with the real-world testing plan offers no direct protection against administrative fines under the EU AI Act.

Pro-innovation measures for SMEs, including start-ups, and microenterprises (Recs.139 and 145; Arts.58 and 62)

– The EU AI Act contains innovation-friendly measures with the aim to reduce the regulatory and administrative burden on SMEs, including start-ups and microenterprises, and to avoid fragmentation of participation.

- SMEs, including start-ups, are given priority access to the sandboxes, provided that they: (1) have a registered office or branch in the EU; and (2) fulfil the relevant eligibility conditions and selection criteria (Art.62(1)(a)).
- Sandboxes will be free of charge for SMEs, including start-ups, other than exceptional costs that a national competent authority may fairly and proportionately recover (Art.58(2)(d)).
- Processes and administrative requirements for participating in and exiting a sandbox will be simple, easily intelligible, and clearly communicated (Art.58(2)(g)), noting the requirement for dedicated channels of communication for SMEs (and other deployers and innovators), where appropriate (Art.62(1)(c)).

Context and illustrations



Commentary: AI sandbox developments

On 9 November 2023, Spain [established](#) a controlled test environment in compliance with the EU AI Act (which was in draft at that time) designed to enable participants to implement high-risk AI systems under the EU AI Act with the aim of obtaining guidance on achieving compliance. Interest in Spain's pilot sandbox has proved strong. In April 2025, 12 high-risk AI projects covering biometrics, employment, critical infrastructure, finance, emergency services, healthcare, and industrial machinery were provisionally admitted to the first cohort.

The sandbox will continue until the EU AI Act is fully applicable in Spain (i.e., 2 August 2026). After that, the results are expected to be published in a report containing conclusions and good practices, which will also be shared with the Commission to help shape EU-level guidance and national regulations.

Accordingly, Spain's pilot sandbox could cause a ripple effect for other Member States to either join in or create their own national sandboxes.

While there have been no other formal announcements at the time of writing, progress appears to be underway for certain other Member States. For example, it is understood that the Danish, Lithuanian, and Swedish data protection authorities will establish or participate in a sandbox. Businesses should stay close to in-region regulatory developments on the creation of sandboxes and associated implementing acts to assess whether they could, and want, to participate.

Regulatory sandboxes will also be available outside of the EU: The UK published a [white paper](#) in March 2023 on AI regulation in the UK which included a plan to create a new £2 million regulatory sandbox to enable businesses to test their products.



Example: Testing of remote biometric identification systems

Company Y (a provider) is developing an AI facial recognition system for remote biometric identification systems and wants to test a prototype.

If Company Y decides to test the prototype in a regulatory sandbox, this would have the advantage of enabling Company Y to use personal data collected for other purposes while testing its AI system. For that, Company Y would have to find an adequate sandbox, apply for it, and be chosen by the responsible authority. Company Y would have to adhere to the sandbox plan and the terms and conditions set by the authority, as well as fulfilling the relevant requirements of other applicable EU laws (notably the GDPR). The results of the test would be part of an exit report. The testing would become public as part of the abstract of the annual report of the sandbox authority. Depending on the sandbox plan, the test could be conducted without a time limit.



Alternatively, Company Y might instead consider testing the prototype in real-world conditions. The advantage of this kind of testing is that Company Y would control the testing. However, a market surveillance authority would still have to accept a testing plan Company Y has written, and Company Y would have to comply with further formal requirements such as registering the testing, being established in the EU or having a legal representative in the EU, obtaining informed consent for the use of personal data, and keeping the test under oversight (as well as ensuring any outcomes are reversible). In addition, Company Y would still need to fulfil the relevant requirements of other applicable EU laws (notably the GDPR). This testing is supposed to end after six months.

Commentary: Will sandboxes actually encourage innovation?

Regulatory sandboxes under the EU AI Act are intended to support innovation by allowing developers to test AI systems in a controlled, compliant environment. However, the EU AI Act's approach to sandboxes may not have the desired effect – and may even inhibit (rather than encourage) innovation.

First, the EU AI Act's measures in support of innovation need to be considered in the context of the rest of the EU AI Act, which (in general terms) prioritises the protection of fundamental rights and the minimisation of harm that could arise from AI systems, ahead of innovation.

Second, participation in sandboxes under the EU AI Act is likely to involve navigating complex approval processes and meeting strict eligibility criteria, which may be burdensome – especially for smaller businesses. This may discourage precisely the kind of experimental projects that drive technological breakthroughs.

Third, sandboxes offer only temporary relief from regulatory uncertainty. Once an AI system exits the sandbox, it is fully subject to the broader, and still evolving, regulatory landscape. As a result, the EU AI Act's sandbox system provides only limited protection from regulation.

Finally, innovation often depends on open ecosystems and informal, iterative development. By isolating participants within controlled environments, sandboxes may limit collaboration and slow down feedback cycles. While well-meaning, there remains a risk that the EU AI Act's sandboxes could become bureaucratic bottlenecks rather than catalysts for progress, favouring compliance over creativity.

Chapter 17

Regulatory framework

Executive summary

The EU AI Act creates a regulatory framework that is designed to enhance coordination at the national level, build capabilities at the EU level, and integrate AI stakeholders. Whether it will achieve these goals remains to be seen.

Relevant sections of the EU AI Act

This Chapter focuses on Chapter VII (**Articles 64 – 70**) of the EU AI Act and the regulatory framework established at the EU level to ensure effective implementation. These governance provisions will apply from 2 August 2025.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **AI Board** (Arts.6(5) and 66) – The European Artificial Intelligence Board is a representative advisory board established to facilitate consistent and effective implementation of the EU AI Act.
- **AI Office** (Arts.3(47) and 64) – The Commission’s function of contributing to the implementation, monitoring and supervision of AI systems and GPAI models, and AI governance. The EU AI Act states that references to the AI Office should therefore be construed as references to the Commission.
- **Advisory Forum** (Art.67) – The Advisory Forum is established to provide technical expertise and advice to the AI Office and the AI Board.

Analysis

Governance bodies – The EU AI Act creates a system of governance at both the EU level and at the national level, including (i) a function within the Commission (the AI Office), (ii) the AI Board composed of EU and Member State official representatives, (iii) an Advisory Forum involving various stakeholders, (iv) an independent Scientific Panel, and (v) a total of 54 designated national competent authorities, i.e., one notifying authority and one market surveillance authority per Member State.

- **AI Office** (Arts.3(47) and 64) – The AI Office was established on 24 January 2024 as a function of the Commission’s Directorate General for Communication Networks, Content and Technology. The AI Office has the powers to regulate GPAI models under Art.75(1) and the Market Surveillance Regulation. The EU AI Act states that references to the AI Office should therefore be construed as references to the Commission. Member States are required to facilitate the tasks allocated to the AI Office (Art.64(1)). The AI Office is responsible for “contributing to the implementation, monitoring and supervision of AI systems and general-purpose AI models” (Art. 3(47)). Its various duties and enforcement powers are laid down in Commission Decision C/2024/1459 and in various provisions of the EU AI Act (Arts.3(47) and 64). These include various tasks aimed at setting up the regulatory framework under the EU AI Act, adopting secondary legislation envisaged by the EU AI Act, enforcing the EU AI Act in relation to GPAI models (with the power to impose fines up to 3% of a company’s annual total worldwide turnover in the preceding financial year or EUR 15 million, whichever is higher – see Chapter 22 for further information on penalties), and ex-post evaluation of the EU AI Act.
- **AI Board** – The AI Board is a representative advisory board established to facilitate consistent and effective implementation of the EU AI Act (Arts.65 and 66). The AI Board is composed of one representative per Member State. The AI Board also includes non-voting observers from the EDPS and the AI Office. The Board adopts rules by a two-thirds majority and is responsible for several tasks, including the issuing of advice, recommendations, and written opinions.

□ **Advisory Forum** – The Advisory Forum is established to provide technical expertise and advice to the AI Office and the AI Board. It will consist of a balanced mix of commercial and non-commercial stakeholders from industry, start-ups, SMEs, civil society, and academia. The Commission will appoint members to the Forum with recognised AI expertise, for a period of up to four years. The Fundamental Rights Agency, the EU Agency for Cybersecurity (ENISA), the European Committee for Standardisation (CEN), the European Committee for Electrotechnical Standardisation (CENELEC), and the European Telecommunications Standards Institute (ETSI) are permanent members of the Forum. The Forum operates under its own internal governance rules, with two elected co-chairs, and meets at least twice a year. It will publish an annual report of its activities (Art.67).

□ **Scientific Panel** – The Scientific Panel will comprise members chosen by the Commission based on their scientific or technical expertise in the field of AI and their independence from any AI system providers. The Scientific Panel will provide impartial advice to the AI Office and support it by helping develop evaluation tools, advising on the classification of AI models and alerting the AI Office to potential risks. EU Member States may ask experts of the Scientific Panel to support their enforcement activities under the EU AI Act, at fees to be determined (Art.68).

□ **National competent authorities** – Member States must, by 2 August 2025, establish or designate as national competent authorities at least one notifying authority and one market surveillance authority for the purpose of implementing the EU AI Act. This may be a pre-existing national authority with other competences (e.g., a data protection authority, or telecommunications regulator). Notifying authorities are responsible for regulating the conformity assessment bodies, which, in turn, demonstrate compliance with high-risk AI rules. Market surveillance authorities are the main enforcement bodies of the EU AI Act. The rules governing national enforcement shall be laid down by Member States but shall include the power to impose fines and take other enforcement measures, such as warnings and non-monetary measures. See Chapter 22 for further information on penalties (Art.70).

□ **EDPS** – The EDPS acts as the competent authority for the supervision of EU institutions, bodies, offices, or agencies where they fall within the scope of the EU AI Act, and may impose fines on EU institutions, bodies, offices, and agencies where they do not comply with the EU AI Act (Art.70(9)).

Who's in charge? – Despite the EU AI Act's call for harmonisation, governance and enforcement may still vary across the EU.

The EU AI Act provides for a two-tiered regulatory framework, with enforcement at both the Member State level and EU-wide level. The Commission – in the form of the AI Office – sits at the heart of this framework, with the power to enforce the EU AI Act's rules regarding the provision of GPAI models (unlike the position under the GDPR, where enforcement is solely at the national level). The AI Office is responsible for providing guidelines, including on consistent enforcement of the EU AI Act (Art.96(1)(e)) (see Chapter 20) and providing coordination support for joint investigations (Art.74(11)). The Board also has a role in contributing to coordination among national competent authorities (Art.66(a)). However, notwithstanding these mechanisms to encourage and facilitate cooperation, the EU AI Act does not provide an explicit power to enable the AI Office to *compel* national competent authorities to coordinate or to adopt uniform interpretations of the EU AI Act. As a result, there remains a risk that different national competent authorities will adopt different interpretations of the EU AI Act.

Context and illustrations



The AI Office

The AI Office, established within the Commission in January 2024, is empowered to oversee the enforcement and implementation within Member States of the EU AI Act.

In its role as an enforcement body, the AI Office will devise codes of practice and guidance for compliant practices under the EU AI Act, as well as investigate possible infringements and publish findings.

It is hoped that the guidance to be issued by the AI Office will provide clarity on the practical application of the EU AI Act to real-world scenarios.

In addition to its role as an enforcement body, a key purpose of the AI Office is to engage in international cooperation and discussion of issues in the AI sector, in recognition of the globalised market for AI products and services and the need for international alignment on ethical AI practices.



The importance of coherent decision-making and cooperation

Each Member State must establish at least one notifying authority to manage compliance and certification processes, and one market surveillance authority to verify that AI systems comply with EU harmonisation legislation (see Chapter 9). These activities and tasks may be performed by one or more authorities, depending on each Member State's organisational needs.

With multiple national competent authorities, their ability to cooperate and make consistent decisions is crucial.

For example, notifying authorities designate and oversee notified bodies, which certify high-risk AI systems. However, there remains a risk that different interpretations and approaches among national authorities could lead to divergent criteria for certification. This creates a risk of "conformity shopping", where providers seek certification in jurisdictions with more lenient practices, undermining the uniform application of EU law.



Divergent approaches in market surveillance authority designation

Market surveillance authorities are responsible for enforcing compliance and investigating non-compliance. However, disparities in their resources and expertise could lead to uneven enforcement across Member States.

EU data protection authorities have advocated that they should be designated as market surveillance authorities. However, Member States have so far followed divergent approaches, which include:

- Denmark, designating the pre-existing Danish Agency for Digital Government.
- Finland, which plans to designate the Finnish Transport and Communications Agency, also responsible for enforcing the DSA.
- Malta, designating the Malta Digital Innovation Authority and the Information Data Protection Commission to jointly serve as market surveillance authorities.
- Spain, which established the new Spanish Agency for the Supervision of Artificial Intelligence, the first of its kind in the EU, and designated it as the market surveillance authority in Spain.

Chapter 18

EU database for high-risk AI systems

Executive summary

With the aim of increasing transparency, oversight, and scrutiny of AI systems, the EU AI Act requires providers of certain high-risk AI systems to register themselves and their systems in an EU database (the “Database”).

As part of that registration, providers are required to provide information about the relevant AI systems, how they work, and what inputs they use. A more limited form of registration applies to public authorities and EU institutions.

These registration obligations overlap with other compliance obligations under the EU AI Act. Businesses therefore need to ensure that they approach registration in a coordinated manner that is consistent with their other compliance efforts.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.49 and 71** and **Annexes VIII and IX** of the EU AI Act (i.e., the key provisions and annexes governing the Database, and the obligation to register).

This Chapter also addresses **Arts.6 and 16** and **Annex III**, which concern the classification of high-risk AI systems, and certain compliance obligations of the providers of such systems, as well as **Art.60** on the testing of high-risk AI systems in real-world conditions, to the extent that those provisions and annexes are relevant to the Database.

Some of the requirements relating to high-risk AI systems are effective 2 August 2026. See Chapter 23 (Commencement and timing) for more information.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Database** (Rec.131; Arts.71 and 49) – An EU database for high-risk AI systems to be set up and maintained by the Commission in collaboration with the Member States, containing certain information (see Annexes VIII and IX) on high-risk AI systems under Art.6(2), Annex III, and not-high-risk AI systems under Art.6(3), to be submitted by the respective provider or, in case of public authorities, the deployer.
- **AI system** (Rec.12; Art.3(1)) – A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.
- **Provider** (Art.3(3)) – A person or entity that develops an AI system or a GPAI model, or that has an AI system or a GPAI model developed, and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
- **Deployer** (Rec.13; Art.3(4)) – Any organisation using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Compliance obligations – The requirement to register high-risk AI systems in the Database is a key obligation for providers (and, in the case of public authorities, deployers) of those systems.

Effective 2 August 2026, businesses that provide high-risk AI systems need to compile the required information and prepare the corresponding entries for the Database by the start of enforcement. Given the level of detail that needs to be submitted to the Database, each business must carefully coordinate these submissions with the rest of its EU AI Act compliance program.

Three key categories of AI systems – The EU AI Act sets out three key categories of AI systems for which the relevant providers (or, in case of public authorities, deployers) must submit certain information to the Database, depending on the respective category:

- **Category 1** – This category covers high-risk AI systems that fall within Art.6(2) and are listed in Annex III. Businesses that act as providers of high-risk AI systems in Category 1 need to register those AI systems in the Database (Arts.71(1), (2), and 49(1)). Registrations for high-risk AI Systems in Category 1 need to include the information required under Section A of Annex VIII (e.g., a “*basic and concise description of the information used by the system (data, inputs) and its operating logic*”).

If a provider intends to test a high-risk AI system listed in Annex III in real-world conditions (outside of an AI regulatory sandbox), certain information specified in Annex IX must be included in the Database registration by the provider (or prospective provider) (Art.60(4)(c)) with the information being accessible only to market surveillance authorities and the Commission without the consent of the provider or prospective provider (Art.71(4)).

- **Category 2** – This category covers AI systems that fall within Annex III but are not considered to be high-risk on the grounds set out in Art.6(3). The fact that an AI system falls within Category 2 means that the provider deems that AI system not to be high-risk. Nevertheless, the provider needs to register that AI system in the Database (Arts.71(1), (2), 49(2), and 6(4)).

Registrations for AI systems in Category 2 need to include the information required under Annex VIII, Section B (e.g., a “short summary of the grounds on which the AI system is considered to be not-high-risk in application of the procedure under Article 6(3)”).

- **Category 3** – High-risk AI systems listed in Annex III, which are deployed by (or on behalf of) public authorities and EU institutions (and anyone acting on their behalf) also need to be registered in the Database (Arts.71(3) and 49(3)). Registrations for AI systems in Category 2 need to include the information required under Annex VIII, Section C (e.g., a “*summary of the data protection impact assessment carried out in accordance with Article 35 of [the GDPR...]*”).

AI systems outside the three main categories – Certain AI systems fall outside the three main categories, but are nevertheless subject to registration obligations:

Substantial modifications of high-risk AI systems will also need to be registered in the Database (Rec.131).

High-risk AI systems in the area of critical infrastructure, i.e., AI systems listed in paragraph 2 of Annex III, will only need to be registered at national level (see Art.49(5), Rec.131).

High-risk AI systems in the area of critical infrastructure (see paragraph 2 of Annex III to the EU AI Act) only need to be registered at national level.

Submitting the required information – The information that needs to be submitted, and the party responsible for that submission, is determined according to the category into which the AI system falls (see above). The required information is specified in Annexes VIII and IX.

The required information regarding relevant AI systems must be compiled and submitted to the Database as follows:

- **For Categories 1 and 2** – The required information needs to be submitted to the Database by the provider of the relevant AI system (or, where applicable, by the provider's authorised representative) (Arts.71(2), 49(1), and (2)).
- **For Category 3** – The required information needs to be submitted to the Database by the public authority or EU institution that is the deployer of that AI system (or, where relevant, by anyone acting on their behalf) (Arts.71(3) and 49(3)).
- **Note** that other deployers that are **not** public authorities or EU institutions are entitled to voluntarily provide information to the Database (see Rec.131).

The information to be submitted by providers or deployers varies depending on which of the three categories of AI systems (see above) applies.

Public and non-public sections of the Database – The Database will consist of a public section and a non-public section, with the latter containing more limited information on high-risk AI systems in certain sensitive areas.

Public section of the Database – This section includes any information contained in registrations filed by the relevant provider/deployer, with the exception of certain registrations for high-risk AI systems in certain sensitive areas (see below; Art.71(4)).

- The public section of the Database is accessible free of charge. It is intended to be user-friendly, with the provided information being easily navigable, understandable, and machine-readable (Rec.131; Art.71(4)).

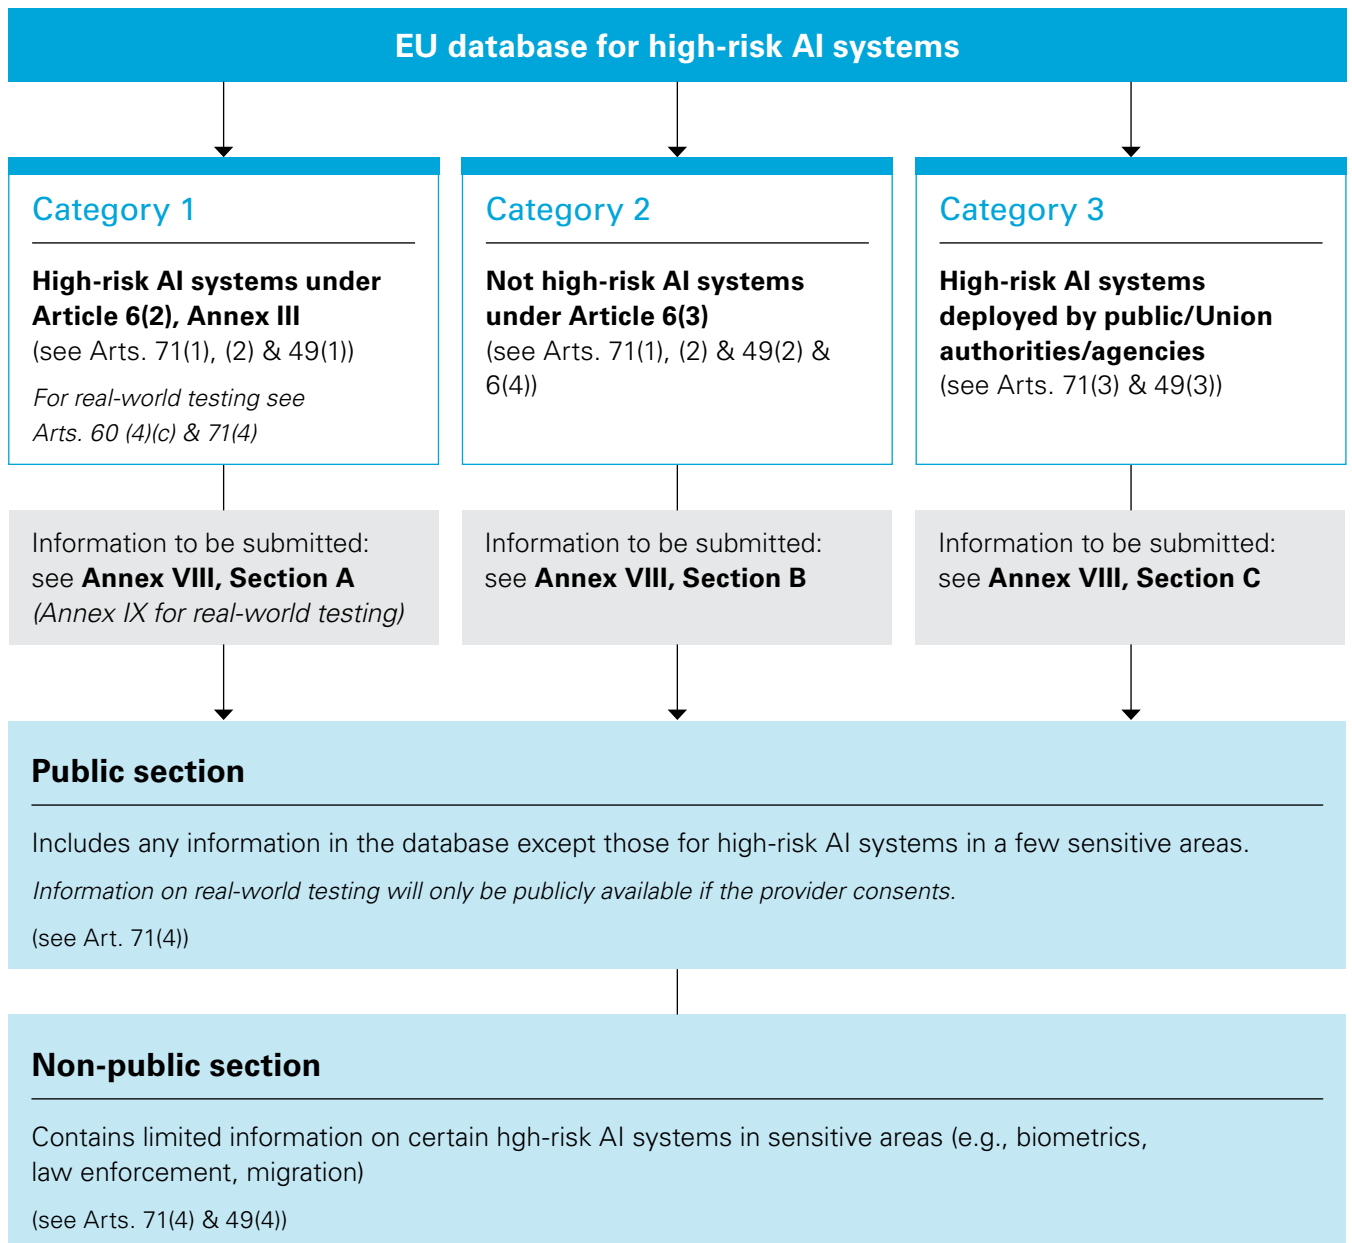
- The information submitted for the testing of high-risk AI systems in real-world conditions will generally be accessible only to market surveillance authorities and the Commission, unless the provider or prospective provider agrees to public availability (Art.71(4)). If the tested high-risk AI system falls into one of the sensitive areas listed below, the registration information will be included in the non-public section of the Database.

Secure non-public section of the Database – This section contains limited information on high-risk AI systems in certain sensitive areas, namely biometrics, law enforcement, migration, asylum, and border control management, including information on real-world testing of those systems (Arts.71(4), 49(4), and 60(4)(c)).

- Access to this section of the Database is strictly limited to the Commission, and to market surveillance authorities with regard to their national section of the Database.

The Commission will play a key role in the establishment and administration of the Database and will also be the controller, for the purposes of EU data protection laws.

The Database will be **established and managed** by the Commission in collaboration with Member States, and the Commission will also act as **the controller**, with respect to the processing of personal data contained in the Database, for the purposes of EU data protection laws (Regulation (EU) 2018/1725; Rec.131, Art.71(1) and (6) of the EU AI Act). **Personal data** should only be included in the Database to the extent necessary for collecting and processing information in accordance with the EU AI Act (Art.71(5)). As stated in Art.71(5), "[t]hat information shall include the names and contact details of natural persons who are responsible for registering the system and have the legal authority to represent the provider or the deployer, as applicable." Additionally, the Commission is tasked with developing functional specifications of the Database and an independent audit report (see Rec.131).



Note: High-risk AI systems in the area of critical infrastructure (see Annex III, Point 2) will only need to be registered at national level (see Art. 49(5), Rec. 131).

Context and illustrations



Example: Registration obligations for AI systems in Category 2

A company develops an AI system that is intended to be used to help determine admission to educational institutions and plans to launch it on the EU market after 2 August 2026. The company (in its role as the provider of the AI system) considers that the AI system is not a high-risk AI system, applying the criteria in Art.6(3) (on the basis that the company does not consider the AI system to pose any significant risk of harm to the health, safety, or fundamental rights of natural persons, since the AI system does not materially influence the outcome of decisions for or against admission). What registration obligations arise for this company?

An AI system that is intended to be used to determine admission to educational institutions risks being classified as high-risk under Art.6(2) and paragraph 3(a) of Annex III. However, if the provider considers that this AI system does not pose a significant risk of harm in accordance with Art.6(3), the AI system falls into Category 2 (as defined above) for the purposes of the Database, and the registration obligations set out in Category 2 (above) will apply, and the provider will need to submit a registration containing the information set out in Section B of Annex VIII. These registration details will be included in the public section of the Database (Arts.49(4) and 71(4)).



Commentary: Significant investment of resources and comprehensive compliance strategy may be required

As the examples demonstrate, a significant investment of resources may be required for businesses to satisfy their registration obligations for the Database. For example, in the case of sophisticated and complex high-risk AI systems in Category 1, it may be challenging for providers to prepare a “*basic and concise description of the information used by the system (data, inputs) and its operating logic*” (see paragraph 6 of Section A of Annex VIII).

Moreover, the provision of some of the information to be included in the Database goes hand-in-hand with the fulfilment of other compliance obligations under the EU AI Act. For AI systems in Category 2, paragraphs 6 and 7 of Section B of Annex VIII require the provider to explain why it concluded that the AI system was not high-risk – a task that may require disclosure of legal advice that would otherwise have been privileged.

Similarly, for an AI system in Category 3, paragraph 5 of Section C of Annex VIII requires the provider to provide a summary of the key findings of its Data Protection Impact Assessment under Art.35 GDPR (where applicable) even though the GDPR itself does not require publication of this information.

Providers should not consider their registration obligations in isolation from their other compliance obligations under the EU AI Act but should instead ensure that they approach registration in a coordinated manner that is consistent with their compliance program as a whole.



Example: Registration obligations for high-risk AI systems in sensitive areas

Company A deploys an AI system, provided by Company B, for the remote biometric identification of persons in the context of border control management. Company A places the AI system on the EU market under its own name on 1 January 2026 and fundamentally overhauls the design of the AI system every six months. Company A does not deploy the system on behalf of EU or other public authorities or institutions. What registration obligations arise for Company A?

Although Company A has deployed the AI system developed by Company B, Company A is still classified as the “*provider*” if it places the AI system on the EU market under its own name (Art.3(3)). The AI system is high-risk (paragraph 1(a) of Annex III). Because Company A does not deploy the AI system on behalf of public authorities and EU institutions, it falls into Category 1. Therefore, Company A must provide the registration information listed in Section A of Annex VIII.

Because the AI system is a biometric system under paragraph 1 of Annex III and was developed for the purpose of border control management, it falls into the non-public section of the Database, (Art.49(4)(a)) (and the information in paragraphs 6, 8, and 9 of Section A of Annex VIII will not be included in the Database at all).

Chapter 19

Monitoring and oversight

Executive summary

The EU AI Act requires providers to conduct ongoing monitoring, oversight, and reporting throughout the lifecycle of high-risk AI systems.

The EU AI Act also contains enforcement provisions that empower various national and EU authorities to conduct market surveillance and impose corrective actions where necessary.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.72 – 94** of the EU AI Act, which lays down the rules for post-market monitoring, information-sharing, and market surveillance. The Chapter includes insights on monitoring and reporting obligations for providers, as well as on enforcement mechanisms available to competent national authorities and the AI Office. This Chapter also covers oversight of GPAI model providers.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Post-market monitoring system** (Art.3(25)) – All activities carried out by providers of AI systems to collect and review experience gained from the use of AI systems they place on the market or put into service for the purpose of identifying any need to immediately apply any necessary corrective or preventive actions.
- **Market surveillance authority** (Art.3(26)) – The national authority carrying out the activities and taking the measures pursuant to Regulation (EU) 2019/1020.
- **AI Office** (Arts.3(47) and 64) – The Commission’s function of contributing to the implementation, monitoring, and supervision of AI systems and GPAI models, as well as AI governance. The EU AI Act states that references to the AI Office should therefore be construed as references to the Commission.
- **National competent authority** (Art.3(48)) – A notifying authority or a market surveillance authority. As regards AI systems put into service or used by EU institutions, agencies, offices, and bodies, references to national competent authorities or market surveillance authorities are construed as references to the EDPS.
- **Serious incident** (Art.3(49)) – An incident or malfunctioning of an AI system that directly or indirectly leads to any of the following:
 - (a) The death of a person, or serious harm to a person’s health.
 - (b) A serious and irreversible disruption of the management or operation of critical infrastructure.
 - (c) The infringement of obligations under EU law intended to protect fundamental rights.
 - (d) Serious harm to property or the environment.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Post-market monitoring and reporting of serious incidents (Arts.72 and 73) – Arts.72 and 73 impose monitoring and reporting obligations on providers of high-risk AI systems:

- **Monitoring obligations for providers (Art.72)** – Providers of high-risk AI systems must establish an appropriate post-market monitoring system that actively and systematically collects, documents, and analyses relevant data. The monitoring system must be proportionate to the nature and risks of the AI system. The monitoring system should allow the

provider to evaluate compliance with the requirements for high-risk AI systems in Arts.8 – 15 (see Chapter 7). The monitoring system must be based on the post-market monitoring plan that providers must make available *before* placing a high-risk AI system on the market. The Commission is expected to provide guidelines for the implementation of a post-market monitoring plan by 2 February 2026.

- **Reporting obligations for providers relating to serious incidents (Art.73)** – Providers of high-risk AI systems must immediately *report any serious incident to the market surveillance authority* of the Member State in which the incident occurred (depending on the circumstances and the nature of the incident, it appears that this may require reporting in many – or possibly all – Member States). The EU AI Act contains different time limits for notification based on the seriousness of the incident, the longest of which is 15 days. For example, if an AI system has malfunctioned, and this has led to a person's death, the incident must be reported within ten days after the provider becomes aware of it. The notified market surveillance authorities will then inform the relevant national public authorities, who will then notify the Commission if the incident is serious. High-risk AI systems that are referred to in Annex III and that fall in the scope of either the Medical Device Regulation or the In Vitro Diagnostics Regulation must also fulfil equivalent reporting obligations. To facilitate compliance with the specific reporting obligations, the Commission is expected to issue guidance by 2 August 2025.

Enforcement (Arts.74 – 84) – National competent authorities and the Commission control compliance of AI systems with the EU AI Act's requirements through market surveillance. In cases of non-compliance, the relevant national competent authorities also have the power to impose corrective measures on businesses as well as to restrict or prohibit an AI system:

- **Market surveillance** – Art.74 refers to the Market Surveillance Regulation (EU) 2019/1020. The Market Surveillance Regulation provides market surveillance authorities with the right to request information. It also

obliges providers to inform the market surveillance authorities about risks associated with an AI system, cooperate with the market surveillance authorities (Art.4(3) Market Surveillance Regulation), and assist with the investigation and enforcement capabilities of the market surveillance authorities (Arts.14 – 16 Market Surveillance Regulation). Market surveillance authorities must annually report any information relating to the application of EU law on competition rules to the Commission as well as to national competition authorities (Art.74(2)). Notably, the market surveillance procedures set out in Arts.79 – 83 (discussed below in this section) **do not apply to AI systems covered by harmonisation legislation** listed in Section A of Annex I (see Chapter 7) if equivalent procedures already exist under such legislation (Art.74(3)).

Member States must establish or designate and publicise competent market surveillance authorities by 2 August 2025. For high-risk AI systems used by financial institutions, the relevant national financial supervisory authority will act as the market surveillance authority (Art.74(6)). However, Member States have the option to designate a different authority for market surveillance of financial institutions as long as they ensure appropriate coordination (Art.74(3)). Additionally, for high-risk AI systems used in law enforcement, border management, and justice and democracy, Member States must designate the competent data protection authority as the market surveillance authority (Art.74(8)). For high-risk AI systems used in EU institutions, the EDPS will act as the market surveillance authority (Art.74(9)).

- **Mutual assistance** – Where an AI system is based on a GPAI model, and both the system and the model are developed by the same provider, the AI Office has the power to supervise compliance (Art.75(1)). It is unclear how this will work in practice – in particular, it is unclear whether the AI Office will always take on responsibility for supervising compliance in such cases, or whether it will only do so on a case-by-case basis. Market surveillance authorities are required to cooperate with the AI Office and keep the AI Board informed (Art.75(2)).

□ **Supervision of real-world testing** – Where real-world testing is carried out (see Chapter 16), market surveillance authorities have the power to supervise such testing. Providers must notify market surveillance authorities before starting such testing, and may only proceed if authorisation is granted, especially when the testing could impact fundamental rights or safety. Market surveillance authorities can require changes, suspend, or terminate the testing if the requirements of Arts.60 – 61 (see Chapter 16) are not met (Art.76).

□ **Protecting fundamental rights** – Member States are required to designate national authorities to oversee the respect of fundamental rights, including the right to non-discrimination, in the context of high-risk AI systems. Member States must publish a list of those authorities. The national competent authorities responsible for protecting fundamental rights will also have power to request and access relevant documentation in relation to high-risk AI systems (Art.77).

□ **Confidentiality** – All authorities involved in enforcement are required to protect the confidentiality of sensitive information, including trade secrets and personal data, obtained during their activities. However, this confidentiality must be balanced with transparency, especially when disclosure is necessary to protect public interest, health, safety, or fundamental rights (Art.78).

□ **AI systems presenting a risk** – An AI system will be a “*product presenting a risk*” (see Art.3 of the Market Surveillance Regulation) if it represents a risk to the health or safety, or to fundamental rights, of any person. If an AI system is found to be non-compliant, national competent authorities can require corrective actions, withdraw the AI system from the market, and must inform the Commission and other Member States if the risk is not limited to its own Member State (Art.79). If, despite being compliant, a high-risk AI system still presents a risk to the health or safety of persons, or to fundamental rights, or to other aspects of public interest protection, the market surveillance authority must ensure that the risk is eliminated before the system is launched. The authority must require the

operator to take corrective actions to ensure that the system no longer presents the risk when placed on the market. The Commission has the power to decide whether additional measures are necessary (Art.82).

□ **Challenges to non-high-risk classifications** – Providers of AI systems that fall within Annex III can claim an exemption from high-risk status, under Art.6(3) (see Chapter 6). But national competent authorities can challenge that exemption if they believe that the relevant AI system actually poses high risks. If an authority finds that the system should be classified as high-risk, it can require the provider to reclassify and comply with the obligations for high-risk AI systems (Art.80).

□ **EU-level oversight** – Art.81 provides for an EU safeguard procedure, in which a Member State or the Commission can object to an AI-related measure taken by another Member State. The Commission will then review the measure and, if necessary, require all affected Member States to take appropriate measures (e.g., removing the AI system from their market).

□ **Formal non-compliance** – Where a high-risk AI system is non-compliant (e.g., where it is not correctly registered in the EU database, or the EU declaration of conformity has not been drawn up correctly), market surveillance authorities can ask the relevant provider to fix the respective issue. If non-compliance continues, the authority can take further action, such as restricting or prohibiting the high-risk AI system from entering the market or recalling or withdrawing it from the market (Art.83).

□ **AI testing support structures** – The Commission is empowered to establish EU-level support structures to facilitate the testing and experimentation of AI systems (in line with Art.21(6) of the Market Surveillance Regulation).

Remedies (Arts.85 – 87) – The EU AI Act provides for various legal remedies to address non-compliance and allows individuals to file complaints with market surveillance authorities:

- Any person or organisation who believes that there has been an infringement of the EU AI Act can submit a complaint to the relevant market surveillance authority (Art.85).
- Anyone who is subject to a decision made using a high-risk AI system listed in Annex III has the right to a clear explanation of how the system was involved in the decision-making procedure if they consider the decision to have a significant adverse impact on their health, safety, or fundamental rights (Art.86). This right appears to operate in parallel to the broadly similar right in Art.15(1)(h) GDPR.
- The Whistleblower Directive (EU) 2019/1937 (which sets out rules for the protection of reporting persons – i.e., whistleblowers) applies to the reporting of infringements of the EU AI Act (Art.87). This effectively requires Member States to apply whistleblower protections (e.g., confidential reporting channels and protection from retaliation) to reporting of EU AI Act infringements.

Supervising providers of GPAI models (Arts.88 – 94) – The EU AI Act sets out oversight rules regarding the supervision of GPAI models:

- **Enforcement and delegation** – The Commission has exclusive powers to enforce the rules regarding GPAI Models (see Chapters 13 and 14), and those powers are delegated to the AI Office (Art.88(1)). In addition, market surveillance authorities can ask the Commission to exercise its powers, where necessary and proportionate, to help those authorities fulfil their tasks under the EU AI Act (Art.88(2)).
- **Monitoring of GPAI models** – The AI Office has a broad power to monitor compliance in relation to GPAI models (Art.89(1)). Companies that integrate GPAI models into their AI systems (i.e., downstream providers pursuant to Art.3(68)) have the right to lodge a complaint alleging that the GPAI model in question is non-compliant the EU AI Act (Art.89(2)).

- **Alerts concerning systemic risks** – The Scientific Panel (see Chapter 17) can alert the AI Office if it suspects a GPAI model poses a significant risk at the EU level or fulfils the conditions for systemic risk (see Chapter 12) (Art.90(1)). The Commission, after informing the AI Board, may exercise the powers noted above, for the purpose of assessing the matter (Art.90(2)).
- **Power to request information** – In order to assess compliance, the AI Office has a broad power to enter into discussions with, or request documentation or additional information from, the provider of a GPAI model (Art.91(1)). The provider (or its representative if relevant) is obliged to supply the information requested (Art.91(5)).
- **Power to evaluate GPAI models** – The AI Office also has the power to evaluate GPAI models to assess compliance or to investigate systemic risk (or appoint independent experts to do so on its behalf). Additionally, the Commission has the power to request access to the GPAI model, including the source code, which the provider must supply upon request (Art.92).
- **Power to request measures** – The Commission can ask providers to take specific measures to comply with obligations or reduce risks. The Commission can also ask providers to limit the availability of their GPAI model on the market, including withdrawing the model from the market, where necessary and appropriate (Art.93).
- **Procedural rights** – The rights set out in Art.18 of the Market Surveillance Regulation apply to the EU AI Act as well. For example, the Commission's actions must have a clear legal basis, which the Commission must explain; decisions must be promptly communicated to their addressees, including information on legal remedies and their respective time limits; and providers must be given at least ten working days to respond before actions are taken, unless urgent public interest concerns require swifter action (Art.94).

Context and illustrations



Commentary: The risk of overlapping demands for information

Under Art.58(1)(a) and (e) of the GDPR, data protection supervisory authorities have extremely broad powers to require controllers and processors to provide “any information” necessary for the performance of their tasks. In the context of an investigation, this can impose a heavy burden on businesses to provide supervisory authorities with detailed information concerning their data processing activities.

Similarly, Art.67 of the DSA allows the Commission to require providers of VLOPs or VLOSEs to provide information relating to suspected infringements of the DSA.

Likewise, Article 77 of the EU AI Act allows national authorities responsible for protecting fundamental rights to request “any documentation created or maintained under [the EU AI Act]” in relation to certain high-risk AI systems. Because the GDPR, the DSA, and the EU AI Act overlap in a number of circumstances, there is a material risk that businesses that are subject to these laws may face significant requests for information by separate regulators, all in relation to the same incident or matter.

Businesses should proactively ensure they have robust internal documentation and compliance procedures, to enable their relevant legal teams to quickly gather the relevant information regarding their AI systems and respond efficiently to information requests under each regime.



Example: Obligations for providers of high-risk AI systems

Company X develops an AI system designed to assist in the diagnosis of rare diseases and places it on the market free of charge, which makes the company a provider according to Art.3(3). The AI system analyses patient data, including genetic information, to provide diagnostic recommendations and therefore is considered a medical device under Art.2(1) of the MDR. Given the critical nature of its application, the AI system is classified as high-risk under Art.6(1) in conjunction with Annex I, paragraph 11. Company X, as a provider, must now establish and document a post-market monitoring system. This system should collect and analyse data to ensure continuous compliance with high-risk AI system regulations under the EU AI Act. The monitoring must, for instance, include tracking the accuracy of diagnoses, patient outcomes, or any potential adverse effects.

After some time on the market, the AI system fails to work as intended and incorrectly diagnoses a health condition of a patient, leading to the patient receiving the wrong treatment. As a result, the patient suffers from severe adverse health effects. This could be considered a serious incident under Art.3(49)(a).

Company X has been made aware of the malfunction and has obtained information about the patient’s health conditions following the wrong treatment prompted by its AI system’s misdiagnosis. Therefore, Company X must immediately report the incident to the relevant market surveillance authority, no later than 15 days after it has become aware of the serious incident.



National competent authorities

Member States must establish or designate three types of competent authorities, i.e., market surveillance authorities, notifying authorities, and national public authorities, by 2 August 2025.

Member States have significant discretion in determining the structure and design of these authorities. Accordingly, Member States have proposed or designated authorities that take a range of forms. Some Member States (e.g., Spain and Italy) have designated specific AI authorities. Others (e.g., Ireland and Finland) have adopted decentralised models, appointing multiple existing regulators to each address AI compliance issues in their existing respective areas of competence. However, at the time of writing, many Member States have not yet finalised their position.

Businesses should pay close attention to announcements regarding the appointment of national competent authorities in the Member States in which they operate, as this is likely to affect issues such as regulatory oversight and incident reporting.

Chapter 20

Codes of conduct and guidelines

Executive summary

The EU AI Act creates two types of non-binding measures to encourage voluntary compliance: codes of conduct and guidelines.

Codes of conduct promote the voluntary adoption of best practices for non-high-risk AI systems.

Guidelines help to ensure that stakeholders have clear, practical guidance for complying with the EU AI Act.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts.95 and 96** of the EU AI Act – specifically, the creation of voluntary codes of conduct for AI systems and additional guidelines developed by the Commission about the practical implementation of the EU AI Act.

Analysis

Voluntary codes of conduct (Recs.165 – 166; Art.95)

– Art.95 encourages the development of voluntary codes of conduct for non-high-risk AI systems. These voluntary codes of conduct aim to foster voluntary compliance with certain requirements from the EU AI Act, such as risk management and transparency, even when the relevant AI systems are not under a mandatory obligation to fulfil those requirements.

- **Participants** (Art.95(1), (3)) – Providers, deployers, and industry organisations can draw up codes of conduct, with input from interested stakeholders such as civil society and academia. The AI Office and Member States are tasked with encouraging and facilitating the creation of voluntary codes of conduct.
- **Scope** (Rec.165; Art.95(1)) – Voluntary codes of conduct apply to all AI systems that are not deemed to be high-risk AI systems (on the basis that high-risk AI systems are already required to meet a higher bar of compliance, and therefore voluntary codes of conduct would be redundant for high-risk AI systems).
- **Measuring compliance** (Art.95(2)) – Voluntary codes of conduct are intended to contain clear objectives and key performance indicators to measure compliance against those objectives. Such objectives may include (but are not limited to) ethical guidelines, environmental sustainability, inclusiveness, and diversity in the design of an AI system, as well as consideration for vulnerable groups.
- **Other specific considerations** (Art.95(4))
 - The specific interests and needs of SMEs, including start-ups, should also be taken into account for these codes of conduct.

Voluntary codes are in addition to other **codes of practice** under the EU AI Act – specifically, in relation to transparency obligations under Art.50(7) (see Chapter 11), for the obligations of GPAI providers under Art.56 (see Chapter 15).

The process of drafting voluntary codes of conduct will be facilitated by the AI Office and Member States, and is likely to follow a similar approach to the process for creating formal codes of practice under Art.56. That process involves extensive consultations with working groups, including providers and stakeholders. However, for voluntary codes of conduct under Art.95, the providers, deployers, and/or industry organisations drafting each code will ultimately hold the pen.

Guidelines on the implementation of the EU AI Act

(Art.96) – The Commission is required to create guidelines on the practical implementation of the EU AI Act.

Art.96 is non-exhaustive in its scope (so, in principle, the Commission can develop guidelines on any aspect of the EU AI Act if it chooses to do so). However, Art.96 places specific emphasis on the development of guidelines on particular topics, including: requirements for high-risk AI systems under Arts.8 – 15 (see Chapter 7), prohibited practices under Art.5 (see Chapter 5), consistent enforcement of the EU AI Act and other EU laws, and the definition of an “AI system” under Art.3(1) (which the Commission issued in February 2025 – see Chapter 3).

As with other provisions of the EU AI Act, the Commission is required to consider the needs of SMEs, start-ups, and other affected sectors when creating guidelines. The Commission must also consider overarching AI governance measures, such as the generally acknowledged state of the art on AI, as well as harmonised standards and specifications, including technical specifications.

The non-binding character of the codes of conduct and the Commission’s guidelines must be assessed on a case-by-case basis.

In principle, these are non-binding instruments. However, the Commission’s guidelines are likely to reflect the Commission’s interpretation of the law – which is likely to be persuasive to both courts and national competent authorities.

Context and illustrations



Commentary: The practical impact of voluntary codes of conduct

Businesses need to take a cautious approach to drawing up and opting into a voluntary code of conduct. In general terms, voluntary codes of conduct operate as a “best practice” tool and (as their name suggests) do not have an enforcement mechanism. Voluntary codes of conduct may provide businesses with a mechanism to help demonstrate compliance with certain aspects of the EU AI Act. However, there is no guarantee that compliance with a voluntary code of conduct will be universally accepted as evidence that a business is compliant with the EU AI Act. Notwithstanding the fact that voluntary codes of conduct are intended to contain “*clear objectives and key performance indicators to measure the achievement of those objectives*”, it remains possible that the Commission and/or national competent authorities could take enforcement action against a business, even though that business is compliant with the relevant voluntary codes of conduct.

It remains to be seen how many voluntary codes of conduct will be created under the EU AI Act. Given the significant resources that will be needed to create a voluntary code of conduct in terms of research, coordination of stakeholders, negotiation, and drafting, it is likely that this will be a slow process.



Commentary: The practical impact of Commission guidelines

Businesses should take note of the Commission’s guidelines because, as noted above, those guidelines typically reflect the Commission’s approach to enforcement and are likely to be persuasive to courts and national competent authorities. As a result, it is advisable for businesses to build the positions set out in the Commission’s guidelines into their compliance programs (or have well-reasoned justifications for taking a different approach).

Guidelines issued by the Commission in other contexts are often designed to answer practical questions concerning the implementation of the law, and are often drafted from a regulatory enforcement perspective rather than a business compliance perspective. Nevertheless, the principles and examples set out in such guidelines should help businesses to craft their compliance programs.

Chapter 21

Implementing regulations

Executive summary

The EU AI Act empowers the Commission to adopt delegated acts in several areas. In doing so, the Commission is subject to the European Parliament and the Council's control and assisted by a Member States committee.

Relevant sections of the EU AI Act

The provisions governing the implementation of regulations are laid down in Chapter XI under the title "*Delegation of power and committee procedure*".

Art.97 sets out rules for the exercise of the delegation, whereas **Art.98** refers to Regulation (EU) 182/2011 laying down the rules and general principles concerning mechanisms for control by Member States.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Conformity assessment** (Art.3(20)) – The process of demonstrating whether the requirements set out in Chapter III, Section 2, relating to a high-risk AI system have been fulfilled.
- **GPAI model** (Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.

A full list of defined terms can be found in the [Glossary](#).

Analysis

The EU AI Act empowers the Commission to adopt, where appropriate, implementing acts to ensure uniform application of the EU AI Act, or delegated acts to update or complement the lists in the Annexes to the EU AI Act. This power is subject to certain limitations laid out in Art.97 as well as continuous control:

- **What is a delegated act?** A delegated act is a legal tool that empowers the Commission to amend or supplement certain elements of the EU AI Act, such as updating technical criteria or thresholds. The aim of these delegated acts is to enable the EU AI Act to be kept up-to-date amid rapid technological advancements, without requiring a wholesale legislative procedure for each such change.
- **Temporal limitation and revocation of power** (Art.97(2) and (3)) – The delegation of powers to the Commission is initially limited to five years, starting from 1 August 2024. It is tacitly extended for periods of an identical duration if not opposed by the European Parliament or the Council. Both institutions are entitled to revoke it at any time. This shall not affect the validity of delegated acts already in force.
- **Consultation of experts** (Art.97(4)) – The Commission is required to consult experts designated by each Member State before adopting a delegated act.
- **Notification and objection** (Art.97(5) and (6)) – Once it adopts a delegated act, the Commission is required to notify that delegated act simultaneously to the European Parliament and to the Council. Both institutions can object if they wish. If they do not, the act enters into force within a period of three months of notification.
- **Committee procedure** (Art.98) – The Commission will be assisted by a special committee formed by Member States. Under the procedure set out in Art.5 of Regulation (EU) 182/2011, the committee votes on the relevant measure, and the Commission can only adopt it if a qualified majority is in favour – providing a degree of oversight by Member States.

Relevant sections – The areas in which the Commission may adopt delegated acts are the following:

- **Exemptions from high-risk rules and use-cases for high-risk AI systems** (Art.6(6) and (7); Art.7(1) and (3)) – The Commission is empowered to adopt delegated acts amending the criteria under which an AI system is considered high-risk, by amending the list of use-cases that are deemed to be high-risk, set out in Annex III (see Chapter 6).
- **Technical documentation requirements** (Art.11(3)) – The Commission can adopt delegated acts as needed, to keep the list of required technical documentation in Annex IV up-to-date (see Chapter 7).
- **Conformity assessment** (Art.43(5) and (6)) – The Commission can adopt delegated acts to amend or update the conformity assessment procedures set out under Annexes VI and VII (see Chapter 10). In addition, the Commission can expand the scope of the conformity assessment procedure to certain high-risk AI systems based on an assessment of the relevant quality management system and technical documentation, with a view to preventing or minimising any risks that AI systems pose to health, safety, and fundamental rights.
- **EU declaration of conformity** (Art.47(5)) – The Commission can adopt delegated acts to update or amend the information that is required to be included in a declaration of conformity set out in Annex V (see Chapter 10).
- **Classification as GPAI models with systemic risk** (Arts.51(3) and 52(4)) – The Commission can adopt delegated acts to update or amend the criteria under which a GPAI model is classified as having systemic risk, including the criteria set out in Annex XIII (see Chapter 12). This notably includes the FLOPs threshold used to calculate systemic risk in certain cases, which is likely to become outdated quickly, and may require regular revision.
- **Obligations for providers of GPAI models** (Art.53(5) and (6)) – In order to enable businesses to satisfy certain technical documentation requirements, the Commission can adopt delegated acts to explain how compliance should be measured and calculated, in line with the criteria in Annex XI (see Chapters 12 and 14).

Context and illustrations



Commentary: Overlap between Art.97 of the EU AI Act and similar EU Laws

Most EU regulations concerning areas that are subject to dynamic technical progress include a mechanism that delegates powers to the Commission, to enable the legislation to be updated.

Provisions similar to Art.97 of the EU AI Act can be found in, for example, Art.87 of the DSA, Art.92 of the GDPR, and Art.45 of the General Product Safety Regulation (EU) 2023/988.

The rules that grant the Commission the power to issue delegated acts under the EU AI Act follow a well-trodden path in EU law. They are designed to allow the Commission to react relatively swiftly to technological change, without ultimately taking away power from the EU's legislative institutions. The European Parliament and the Council can object to the adoption of concrete rules as well as revoke the delegation, providing an ultimate backstop if needed.



Multi-stakeholder influence on implementing regulations

The Commission can expand the categories of high-risk AI systems by adopting a delegated act (i.e., the Commission can change the rules so that some types of AI systems that are currently treated as high-risk will become high-risk AI systems). To do this, the Commission will first have to align with Member State experts and then with a committee of representatives of Member States. Once the Commission has adopted the delegated act, it has to notify the Council or the European Parliament, who then have three months to object (and either of them can extend that period by a further three months). If no objection is received, the delegated act takes effect.

Political stakeholders will be able to influence the scope of additional use-cases of high-risk AI via multiple institutions. Notably, Member States have three channels (experts, Commission, and Council) to influence the content of the delegated act.



Commentary: Impact on businesses

For businesses, it is important to keep in mind that due to delegation of powers in various areas, the relevant provisions underlay dynamic change. This means that EU AI Act compliance programs should not be thought of as being aimed towards a fixed target. Instead, EU AI Act compliance programs will require flexibility and adaptation, as the relevant rules are likely to evolve over time. It is therefore essential for businesses to stay up-to-date with the applicable rules as they develop.

Chapter 22

Penalties

Executive summary

The EU AI Act empowers EEA agencies and national bodies to impose penalties, including fines, on providers of AI systems for breaches of the EU AI Act. The maximum fines under the EU AI Act are substantial – the greater of €35 million or 7% of worldwide turnover.

Relevant sections of the EU AI Act

This Chapter focuses on **Arts. 80, 91, 92, and 99 – 101** of the EU AI Act – which establish the penalties that may apply to individuals, organisations, and EU institutions and bodies acting in breach of the EU AI Act.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **Deployer** (Art.3(4)) – Any organisation using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.
- **Distributor** (Art.3(7)) – Any organisation, other than the provider or the importer, that makes an AI system available on the EEA market.
- **EDPS** (Art.70(9)) – The European Data Protection Supervisor is the data protection supervisory authority for the EU institutions, under Art.52 of Regulation (EU) 2018/1725. It is allocated additional supervisory powers in relation to AI, under the EU AI Act.
- **Importer** (Art.3(6)) – Any organisation located or established in the EEA that places on the market an AI system which bears the name of an entity established outside of the EEA.

- **Operator** (Art.3(8)) – A catch-all term for providers, product manufacturers, deployers, authorised representatives, importers, and distributors.
- **Provider** (Art.3(3)) – Any organisation that develops an AI system or a GPAI model, or that has an AI system or a GPAI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.
- **Undertaking** – The term is used several times in the EU AI Act but is not defined. Case law states that an “*undertaking*” includes “*every entity engaged in an economic activity, regardless of the legal status of the entity and the way in which it is financed*” (*Höfner v. Macrotron* (Case C-41/90)). Penalties based on percentages of turnover under the EU AI Act are determined on the basis of the turnover of the relevant “*undertaking*” (see below in this Chapter). As a result, those penalties are **not** calculated on the basis of the turnover of the entity responsible for the relevant infringement. Instead, they are based on the combined turnover of all entities involved in the relevant economic activity (which in many cases may mean the entire corporate group, and in some cases may include third-party entities).

A full list of defined terms can be found in the [Glossary](#).

Analysis

In certain cases, market surveillance authorities have the power to impose fines on providers:

- **Art.80(4)** – If an AI system is classified by the provider as non-high-risk but is found to be high-risk by a market surveillance authority, the provider is required to bring the AI system into compliance with the regulations regarding high-risk AI systems. In the event of continued non-compliance, the market surveillance authority may issue a fine.

- **Art.80(7)** – If a market surveillance authority determines that a provider has intentionally misclassified their AI system as non-high-risk in order to circumvent the requirements for high-risk AI systems, that market surveillance authority may issue a fine.
- **Art.91(4)** – In the event of incorrect, incomplete, or misleading information in response to a request of information by the Commission, the provider of the relevant GPAI model may face fines.
- **Art.92(4)** – If providers of GPAI models do not comply with a request for access to the GPAI model by the Commission to evaluate its compliance, the provider of the relevant GPAI model may face fines.

National rules on penalties and enforcement measures:

Art.99 and Rec.168 confer responsibility on Member States for laying down national rules on penalties for breaches of the EU AI Act, including both monetary and non-monetary, to encourage compliance and ensure that the EU AI Act is enforceable nationally.

- Penalties introduced by Member States must be effective, proportionate, and dissuasive, and must take into account the interests and economic viability of SMEs, including start-ups.
- Art.99 establishes the following limits on fines that can be levied in respect of breaches of particular provisions of the EU AI Act:
 - An administrative fine of up to €35 million or (for undertakings) 7% of worldwide annual turnover for the preceding financial year (whichever is higher), for breaches of any of the prohibitions on AI practices set out in Art.5 (see Chapter 5).
 - An administrative fine of up to €15 million or (for undertakings) up to 3% of worldwide annual turnover for the preceding financial year (whichever is higher), for breaches of the following obligations related to operators or notified bodies (other than those in Art.5): Arts.16, 21 – 24, and 26 (obligations of providers, importers, distributors, and deployers of AI systems) (see Chapter 8); Arts.31 – 34 (obligations relating to notified bodies) (see Chapter 9); and Art.50 (transparency obligations of providers and deployers) (see Chapter 11).

- An administrative fine of up to €7.5 million or (for undertakings) 1% of worldwide annual turnover for the preceding financial year (whichever is higher), for providing incorrect/incomplete or misleading information in response to a request from a regulator.
- For SMEs, the maximum fine is the lower of the percentage or amount specified above.

Article 100 – Administrative fines on EU institutions, bodies, offices, and agencies:

- The EDPS is empowered to monitor and ensure the protection of personal data and privacy when EU institutions and bodies process personal data in the context of the EU AI Act.
- Art.100 empowers the EDPS to impose administrative fines on EU institutions, bodies, offices, and agencies falling within the scope of the EU AI Act.
- The EDPS is required to have regard to various factors under Art.100 when deciding whether and in what amounts to impose administrative fines for breaches of the EU AI Act, including the nature and gravity of any infringements, the degree of responsibility of the relevant institution, any efforts made by that institution to cooperate with the EDPS and to mitigate damage caused to any affected persons, and any previous infringements by the relevant institution.

Article 101 – Fines for providers of GPAI models:

Art.101 empowers the Commission to impose direct fines on providers of GPAI models. Pursuant to Art.101, the Commission may impose fines on providers of GPAI models of up to 3% of their annual worldwide turnover in the preceding financial year, or €15 million (whichever is higher), if a provider intentionally or negligently:

- Infringes the provisions of the EU AI Act.
- Fails to comply with (i) requests for documents or information made pursuant to Art.91; (ii) measures requested under Art.93; or (iii) the obligation to provide the Commission with access to AI models for the purposes of evaluations made under Art.92.

Context and illustrations



Commentary: Fines determined by reference to worldwide turnover of an undertaking

The fact that the EU AI Act specifies the maximum levels of administrative fines by reference to worldwide turnover of the relevant undertaking is consistent with the EU's approach to penalties under other key elements of the EU's Digital Strategy, including the GDPR, the DSA, and the DMA. The focus on worldwide turnover means that businesses with large global operations will not be protected by the fact that their presence in the EEA is relatively small. In fact, as noted in Chapter 2, even businesses with no presence in the EEA may still be subject to the EU AI Act in some circumstances and may therefore face penalties based on worldwide turnover.

As noted in the main body of this Chapter, the term "*undertaking*" means that worldwide turnover is calculated on the basis of the relevant economic unit, regardless of its legal or corporate form. This means that businesses cannot use the corporate veil to prevent portions of their operations from being included in the calculation of worldwide turnover.

The EU's overall objective in setting the EU AI Act's maximum fines so high (noting that they are 175% of the maximum fines under the GDPR) appears to be to ensure that such fines provide an effective deterrent to infringements of the EU AI Act.



Commentary: SMEs and start-ups

Embedded in the penalty provisions of the EU AI Act are measures designed to ensure that any penalties achieve the EU's aims of being effective, proportionate, and dissuasive, while attempting to provide adequate protection for the particular interests and economic viability of SMEs, including start-ups.

The specific reference to the interests of SMEs and start-ups in the EU AI Act, along with the use of lower maximum penalties for SMEs for some breaches, illustrates the EU's attempt to balance enforcement with fairness. This approach aims to ensure the measures of the EU AI Act are suitably enforced while avoiding unduly harsh sanctions that might stifle innovation and the development of novel AI systems and applications within the EEA.

However, it could be said that the calculation of penalties based on percentages of turnover, with special protections on SMEs, amounts to a penalty on success. A multinational company and a small start-up could commit the same violation – even causing the same harm to the same number of individuals – but the larger company would face disproportionately higher financial consequences simply because it is more commercially successful. This arguably creates a disparity where success, rather than wrongdoing, becomes the basis for punishment.



Commentary: How quickly should businesses expect large penalties to be issued?

Immediate enforcement of the EU AI Act on day one is unlikely for most businesses, due to the sheer complexity of its requirements, the need for extensive compliance infrastructure, and the gradual rollout mechanisms built into the legislation. Both national regulators and the AI Office will need time to issue guidance, set up compliance pathways, and coordinate enforcement strategies. In addition, as noted in Chapter 23, the start of enforcement of the EU AI Act is staggered over several years.

As we have seen with other EU laws (especially those that rely on an element of enforcement by national regulators), it takes time for new regulators to be set up and for them to get to grips with their powers. As a result, where the correct interpretation is unclear, regulators may be hesitant to issue very large penalties for fear of being overturned on appeal (whereas businesses might not risk appealing smaller penalties). That said, this is not a universal certainty, and there are some national authorities that may be willing to immediately adopt aggressive enforcement positions and face the risk of being overturned in court.

Accordingly, although businesses may be afforded some time to adapt to the EU AI Act, such lenience may not be universal across the EU.

Chapter 23

Commencement and timing

Executive summary

The EU AI Act has a complex timeline of enforcement, with different provisions applying from different dates. Most of its provisions will apply from 2 August 2026. However:

- Rules regarding prohibitions apply from 2 February 2025.
- Certain rules regarding GPAI models apply from 2 August 2025.
- Certain rules regarding high-risk AI systems apply from 2 August 2026.
- Certain rules regarding some large-scale IT systems apply from 31 December 2030.

Understanding this timeline is crucial for any business that is seeking to develop an EU AI Act compliance plan. In addition, businesses should be aware that several portions of the EU AI Act are likely to be subject to ongoing revision.

Relevant sections of the EU AI Act

This Chapter focuses on Arts.102 – 113 (i.e., the final Chapter of the EU AI Act). It covers the entry into force and application of the EU AI Act, amendments to other EU legislation, and the evaluation and review of the EU AI Act.

Key defined terms

The key defined terms used in this Chapter are as follows:

- **AI system** (Art.3(1)) – A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.
- **GPAI model** (Art. 3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks, regardless of the way the model is placed on the EU market, and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the EU market.
- **Substantial modification** (Art.3(23)) – A change to an AI system after its placing on the EU market or putting into service which is not foreseen or planned in the initial conformity assessment carried out by the provider and, as a result of which, the compliance of the AI system with the requirements set out in Chapter III, Section 2, of the EU AI Act is affected or results in a modification to the intended purpose for which the AI system has been assessed.

A full list of defined terms can be found in the [Glossary](#).

Analysis

Entry into force and application – The EU AI Act entered into force 20 days after its publication in the Official Journal of the EU, i.e., on 1 August 2024, and will apply from 2 August 2026. However, for certain kinds of AI systems and GPAI models, enforcement dates will differ (Arts.111 and 113):

- The enforcement of the EU AI Act did not start on its entry into force on 1 August 2024, but voluntary compliance is encouraged from that date (Rec.178).
- Any use of AI practices prohibited under the EU AI Act (see Art.5) must cease by 2 February 2025 or face enforcement risk (Art.113(a)).
- For GPAI models placed on the market before 2 August 2025, enforcement begins on 2 August 2027 (Art.111(3)).

- For GPAI models placed on the market on or after 2 August 2025, enforcement begins when the GPAI model is placed on the market (Art.113(b)).
- High-risk AI systems that are placed on the market before 2 Aug 2026, and are not intended for use by public authorities, face no enforcement of the rules on high-risk AI systems (as long as there are no significant changes to those AI systems) (Art.111(2)).
- For high-risk AI systems under Art.6(2) that are placed on the market, or are significantly changed, on or after 2 August 2026, the enforcement of the rules on high-risk AI systems begins when the AI system is placed on the market (Arts.6(2), 111(2), and 113).
- For high-risk AI systems under Art.6(1) that are placed on the market, or are significantly changed, on or after 2 August 2026, the enforcement of the rules on high-risk AI systems begins on 2 August 2027 (Arts.6(1), 111(2), and 113(c)).
- For high-risk AI systems that are placed on the market before 2 August 2026 and are not subject to any significant changes after that date, and are intended for use by public authorities, enforcement begins on 2 Aug 2030 (Art.111(2)).
- For AI systems that form part of large-scale IT systems under Annex X, and that are placed on the market before 2 August 2027, enforcement begins on 31 December 2030 (Art.111(1)).

For a visual timeline, please refer to our EU AI Act enforcement timeline at the end of the Chapter. Note that the dates set out above reflect the earliest point at which enforcement could start in each case. It is not yet certain that the relevant regulators will necessarily be in place, or will start enforcement, on each of these dates.

Amendments to other EU laws – The EU AI Act amends a variety of EU laws whose subject matter is affected by AI regulatory developments.

Regulations and Directives amended by the EU AI Act cover a wide range of subject matter and include:

- Regulations (EC) 300/2008 and (EU) 2018/1139, which set out rules governing certain aspects of civil aviation and establish the European Union Aviation Safety Agency.
- Regulation (EU) 167/2013, which sets out rules governing approval of new agricultural and forestry vehicles.
- Regulation (EU) 168/2013, which sets out rules governing two- or three-wheel vehicles and quadricycles.
- Regulation (EU) 2018/858 and (EU) 2019/2144, which set out rules governing motor vehicles and trailers.
- Regulation (EU) 2018/1139, which sets out rules governing certain aspects of civil aviation and establishes the European Union Aviation Safety Agency.
- Directive (EU) 2016/797, which sets out rules governing rail transport services in the EU.
- Directive (EU) 2020/1828, which sets out rules governing representative actions for the protection of the collective interests.

The amendments concern the integration of AI systems and oblige the Commission to take into account the mandatory requirements for high-risk AI systems (see Chapter 7) when adopting any relevant delegated or implementing acts on the basis of those acts.

Ongoing evaluation and review – The evaluation and review of the EU AI Act is entrusted to the Commission.

The Commission will review the EU AI Act, initially on an annual basis, focusing on the need for amendments to Annex III and the list of prohibited AI practices, and submit its findings to the European Parliament and the Council (Art.112(1)). By August 2028, and every four years thereafter, the Commission will evaluate specific areas such as transparency measures, supervision and governance effectiveness, and the status of national authorities' resources, and submit a report to the European Parliament and the Council (Art.112(2) and (4)). The Commission can propose amendments to the EU AI Act in light of technological advancements, health, safety, and fundamental rights impacts (Art.112(10)).

Additionally, the Commission will assess the AI Office's functioning and the development of energy-efficient AI standards (Art.112(5)) as well as review the impact of voluntary codes of conduct every three years (Art.112(7) – see Chapter 20).

Context and illustrations



Commentary: High-risk AI systems already placed on the EU market

The staggered enforcement timeline of the EU AI Act was necessary to give businesses a chance to achieve compliance with respect to AI systems and GPAI models already in development. However, the timeline is likely to have an impact on how businesses approach compliance. In particular, a GPAI model placed on the EU market on 1 August 2025 faces no enforcement risk for two years, while a GPAI model placed on the EU market on 2 August 2025 faces an immediate enforcement risk (Art.111(3)). Providers of GPAI models are therefore incentivised to place those models on the EU market before that deadline.

Similarly, high-risk AI systems that are placed on the EU market before 2 August 2026 will only face enforcement if they are later subject to “*significant changes*” (Art.111(2)). Again, providers of high-risk AI systems are incentivised to place those systems on the EU market before that deadline (and not significantly change them thereafter).

Rec.177 explains that the concept of “*significant change*” (as used in Art.111(2)) is “*equivalent in substance*” to the term “*substantial modification*”. It is not entirely clear why two separate terms are used.



Example: Significant changes to high-risk AI systems already placed on the EU market

Company X, a technology company based in the EU, has developed a virtual health assistant (VHA). VHA is a high-risk AI system under Art.6(1) and is designed to assist healthcare providers with patient diagnosis and treatment recommendations. Before placing VHA on the EU market, Company X conducts a conformity assessment to ensure that the system complies with all the requirements of the EU AI Act (see Chapters 8, 9, and 10). Company X places VHA on the EU market on 1 January 2025.

Scenario 1: After placing VHA on the EU market, Company X does not make any significant changes to VHA. Therefore, VHA faces no enforcement risk with regard to the rules on high-risk AI systems under the EU AI Act.

Scenario 2: Company X decides to make changes to VHA. Company X alters the diagnostic algorithms and integrates new data sources. The changes were not initially foreseen in the conformity assessment. The changes take effect after 2 August 2027. As a result of the changes, VHA is facing an enforcement risk with regard to the rules on high-risk AI systems under the EU AI Act from the point at which those changes go live (Arts.111(2) and 113(c)).

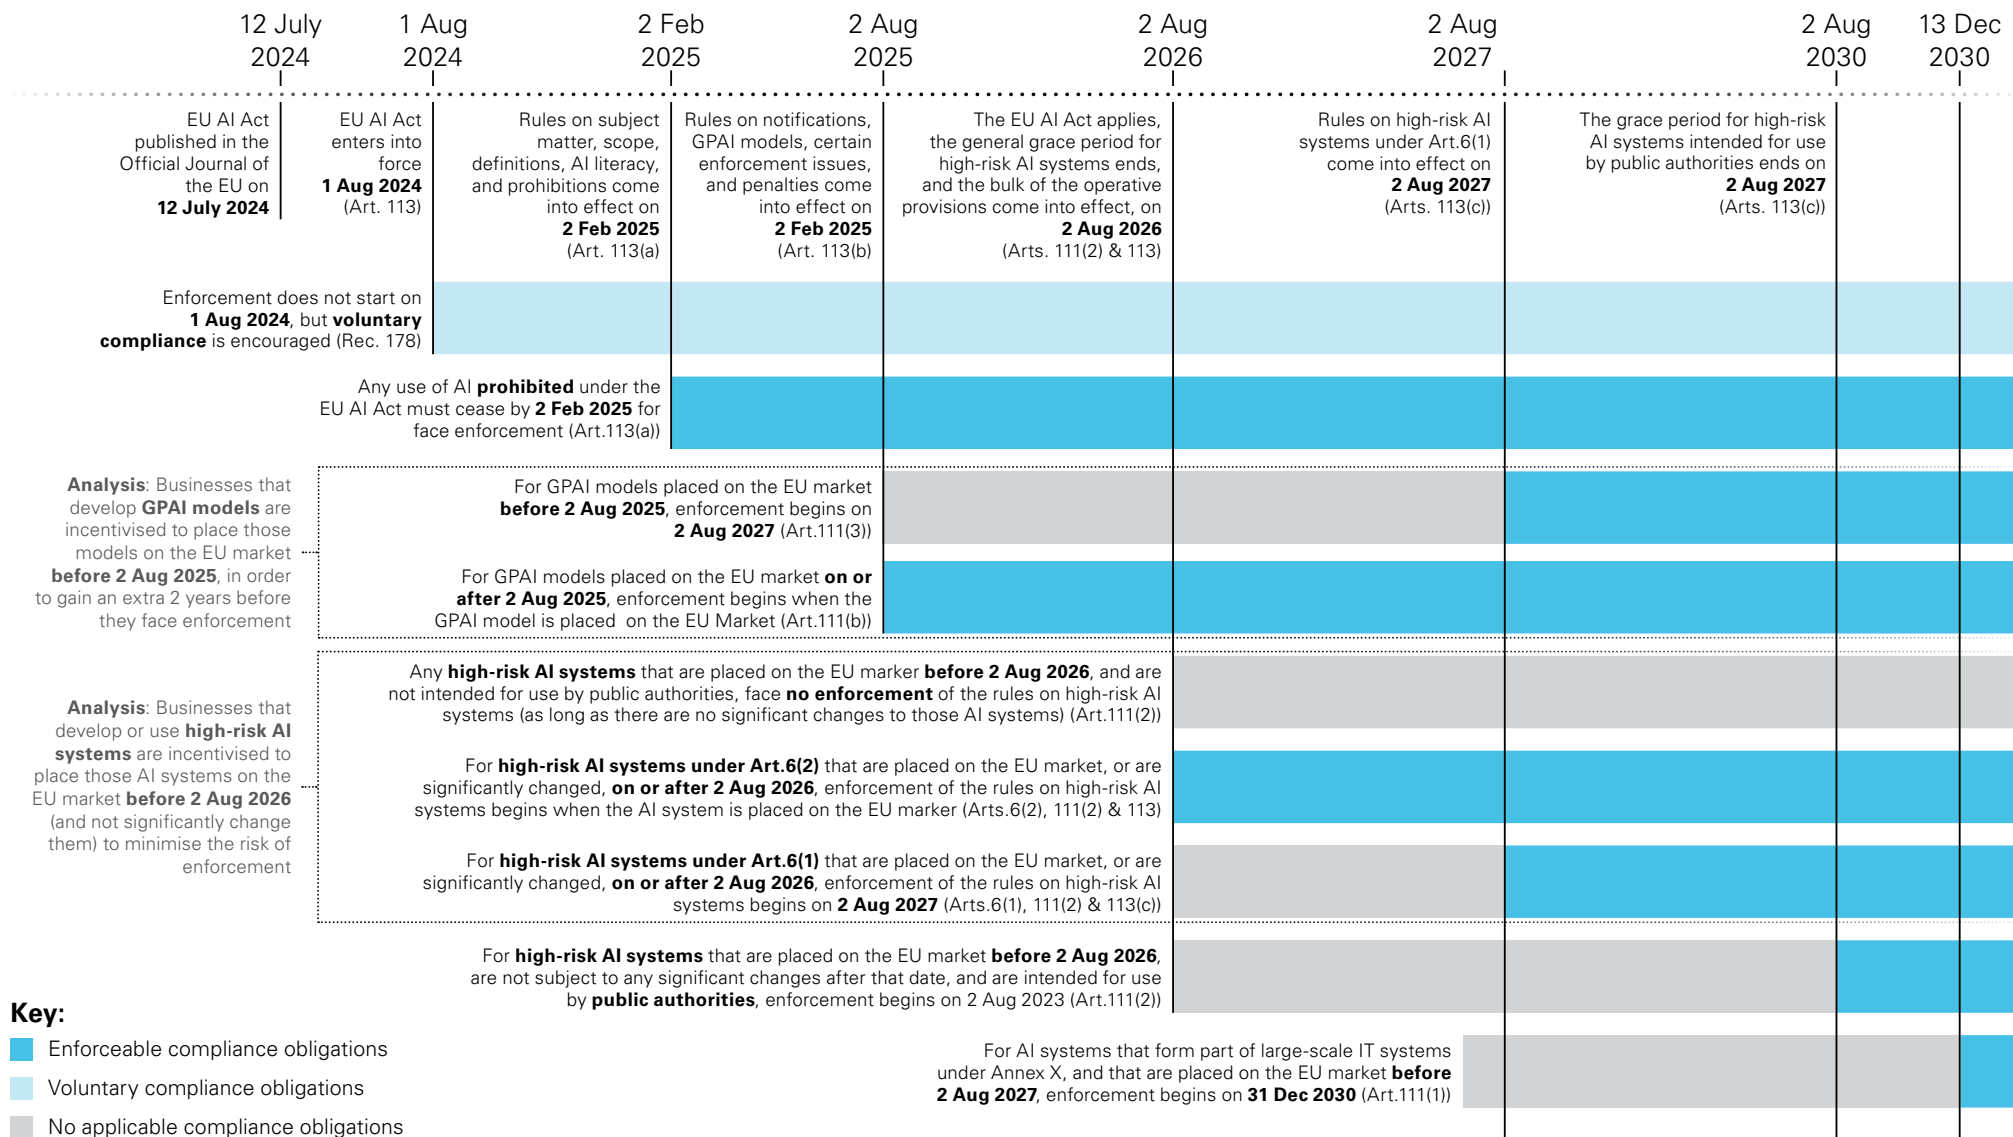


Commentary: What happens if the deadline for enforcement of EU AI Act obligations passes, but the relevant regulators are not yet in place?

If the deadline for enforcement of EU AI Act obligations passes but the relevant market surveillance authorities are not yet fully operational, businesses would face a legal obligation to comply without a clearly functioning enforcement mechanism. However, this does not mean that businesses would automatically have no compliance obligations – in principle, it would be possible for a regulator to later bring enforcement action in respect of alleged non-compliance that took place after the start of enforcement under the EU AI Act, even if it took place before the regulator was in place.

In that scenario, the absence of active regulators would not suspend the effect of the EU AI Act; rather, it would create a transitional vacuum where businesses would likely be expected to act in good faith (noting that voluntary compliance is encouraged even before formal enforcement begins – see Rec.178). In such a scenario, while formal penalties may be delayed, businesses could still face legal uncertainty, reputational risk, and potential liability. However, in that scenario, businesses may be able to point to the lack of a functioning market surveillance authority as a mitigating factor where businesses made compliance choices that could not be checked with the relevant authority at the time.

EU AI Act enforcement timeline



Chapter 24

Overlap between the EU AI Act and other relevant laws and guidelines

Executive summary

The EU AI Act is not the only piece of EU law that regulates the development and use of AI technologies. There are several other laws that impose requirements that affect AI in specific contexts (e.g., data protection, digital services, machinery regulation, medical devices, product liability, and so on). For the most part, the EU AI Act is designed to act in concert with these other laws, but it is important for businesses to be aware of these overlaps.

Laws, draft laws, and guidelines

This Chapter focuses on enacted and forthcoming legislation and soft law related to the EU AI Act. These are instruments applicable within the EEA and the EU.

Specifically, it covers the:

- **General Data Protection Regulation (GDPR)** ([Regulation \(EU\) 2016/679](#)).
- **Digital Services Act (DSA)** ([Regulation \(EU\) 2022/2065](#)).
- **Digital Markets Act (DMA)** ([Regulation \(EU\) 2022/1925](#)).
- **Machinery Regulation** ([Regulation \(EU\) 2023/1230](#)).
- **Medical Devices Regulation (MDR)** ([Regulation \(EU\) 2017/745](#)).
- **Proposed AI Liability Directive (AILD)** ([2022/0303\(COD\)](#)) (now withdrawn).
- **Revised EU Product Liability Directive (rPLD)** ([Directive \(EU\) 2024/2853](#)).
- **Revised Child Sexual Abuse Directive (rCSAD)** ([2024/0035\(COD\)](#)).

- **EDPS First Orientations for generative AI and data protection compliance for EU institutions (the “EDPS Orientations”)** ([link](#)).

For the most part, the EU AI Act is designed to act in concert with these laws and guidelines, but it is important for businesses to be aware of these overlaps. It is also important to be aware that this list is not exhaustive. Because AI can be applied in so many different contexts, it is always possible that there are other applicable laws or guidelines that may affect a business’s ability to lawfully develop or use AI in such contexts.

Analysis

Introduction – With the EU AI Act in the spotlight, it is essential to understand that there are several other laws that will impact the development and use of AI, and more are on the way.

These laws vary in scope and nature, and include EU Regulations that directly impact companies using or developing AI, EU Directives that have to be implemented into national laws of Member States, and guidelines from EU regulatory bodies. This Chapter provides an overview of these laws and guidelines, and emphasises the fact that businesses operating in relevant fields will need to consider these other laws and ensure that their compliance programs for these other laws are coordinated with the equivalent compliance efforts in relation to the EU AI Act.

The GDPR

Key dates and facts:

- **Entered into force:** 25 May 2018.
- **Legal nature:** Binding EU Regulation.

The GDPR and the EU AI Act both regulate automated decision-making, profiling, and the use of personal data. The EU AI Act states that it is “*without prejudice*” to the GDPR. In addition to complying with the EU AI Act, any business that is looking to implement AI technologies that involve the processing of personal data will need to consider Art.22 GDPR – which effectively imposes a prohibition on fully automated decision-making (without human oversight) that results in a legal effect, or similarly significant effect, for individuals. Arts.12 – 14 GDPR impose extensive transparency obligations,

requiring businesses to explain their data processing activities. Businesses that use AI to process personal data will need to ensure that: (i) such processing is accurately explained in their privacy notices; and (ii) their privacy notices are consistent with any transparency information they produce under the EU AI Act. It is also worth noting that a data breach involving an AI system that was being used to process personal data could potentially lead to parallel investigations and penalties under both the EU AI Act and the GDPR.

The DSA

Key dates and facts:

- **Entered into force:** 16 November 2022 (becoming generally applicable on 17 February 2024).
- **Legal nature:** Binding EU Regulation.

The DSA regulates online platforms and recommender systems, including AI-based content moderation. The EU AI Act states that it is “*without prejudice*” to the DSA. As explained in greater detail in Chapter 11, several of the DSA’s requirements with respect to transparency, mitigation measures, and fake content overlap with the requirements of the EU AI Act. Businesses that are subject to the DSA, and that use AI, need to consider their DSA and EU AI Act compliance obligations in parallel, to ensure that their approach to complying with these regimes is consistent.

The DMA

Key dates and facts:

- **Entered into force:** 1 November 2022 (with staggered implementation, becoming generally applicable on 2 May 2023).
- **Legal nature:** Binding EU Regulation.

The DMA and the EU AI Act intersect in several important areas, particularly where large digital platforms deploy AI systems in their services. The DMA imposes obligations on so-called “*gatekeepers*” – very large online platforms that seem to have significant market power – with the aim of regulating competition in the digital market. Where gatekeepers use AI, they may be subject to both the DMA and the EU AI Act in

parallel. Certain services provided by gatekeepers (e.g., personalised ads, content recommendation, ranking systems, and automated moderation) commonly rely on AI and algorithmic decision-making. The DMA imposes obligations regarding transparency in how these systems operate. This overlaps with the EU AI Act’s emphasis on transparency, accountability, and human oversight, especially for high-risk AI systems.

While the DMA focuses on competition issues and platform behaviour, and the EU AI Act governs the technical deployment of AI systems and GPAI models, the DMA and the EU AI Act overlap in imposing obligations to ensure that AI technologies are used in a manner that is deemed to be fair and compliant with fundamental rights and safety standards. As such, companies subject to the DMA will likely also need to comply with EU AI Act requirements when deploying AI systems.

The Machinery Regulation

Key dates and facts:

- **Entered into force:** 19 July 2023 (with staggered implementation, becoming generally applicable on 14 January 2027).
- **Legal nature:** Binding EU Regulation.

The Machinery Regulation establishes health and safety requirements for the design and construction of machinery within the EEA. It replaces the previous Machinery Directive 2006/42/EC. The Machinery Regulation aims to ensure a uniform standard of safety for machinery placed on the EEA market. When machinery incorporates AI components, manufacturers may be subject to both the EU AI Act and the Machinery Regulation. In particular, machinery that includes AI systems must meet the health and safety standards outlined in the Machinery Regulation and also adhere to the risk management and transparency obligations specified in the EU AI Act. Manufacturers may also need to complete conformity assessments that address both machinery safety and AI system compliance. Understanding the interplay between the Machinery Regulation and the EU AI Act is crucial for manufacturers and stakeholders involved in the development and deployment of AI-integrated machinery within the EEA market.

The MDR

Key dates and facts:

- **Entered into force:** 25 May 2017 (with staggered implementation, becoming generally applicable on 26 May 2021).
- **Legal nature:** Binding EU Regulation.

The MDR sets out the legal framework for ensuring the safety, performance, and quality of medical devices marketed in the EEA.

As noted in Chapters 7 – 10, there is significant overlap between the MDR and the EU AI Act. Medical devices utilising AI are generally classified as high-risk AI systems, especially if they also require third-party conformity assessments under the MDR. Both the MDR and the EU AI Act require the implementation of quality management systems. In addition, the MDR requires ongoing monitoring of medical devices after they enter the EEA market. The EU AI Act complements this by requiring continuous monitoring of AI system performance, including mechanisms for reporting incidents related to AI functionality. Manufacturers of medical devices that incorporate AI components need to ensure that they take account of both the requirements of the MDR and the requirements of the EU AI Act.

The AILD (now withdrawn)

Key dates and facts:

- **Proposed:** 28 September 2022.
- **Entered into force:** N/A – withdrawn 11 February 2025.
- **Legal nature:** Binding EU Directive.

The proposed AILD was intended to be applicable to non-contractual fault-based civil law claims brought before Member State courts. It was designed to ease

the burden of proof for alleged victims of harm caused by AI systems, and to make it easier for national courts to obtain evidence about high-risk AI systems that are suspected of having caused damage. However, on 11 February 2025, the Commission [announced](#) that the AILD was being withdrawn, citing a lack of any foreseeable agreement. The Commission also explained that it will “*assess whether another proposal should be tabled or another type of approach should be chosen*”.

The rPLD

Key dates and facts:

- **Proposed:** 28 September 2022.
- **Entered into force:** 9 December 2024.
- **Transposition by:** Two years after entry into force (see Art.22(1) rPLD).
- **Legal nature:** Binding EU Directive.

The rPLD modernises the original 1985 PLD to better address digital technologies, AI, and complex value chains. It introduces strict no-fault liability for a broader range of products, including software and AI systems, whether standalone or embedded. Under the rPLD, software and AI are now explicitly considered “*products*”, and liability extends beyond the point of sale, covering issues like updates, cybersecurity failures, and machine learning. In addition, psychological harm and data loss are now recognised as compensable damages. The burden of proof is eased for claimants in complex cases, with manufacturers required to share relevant information. More actors in the AI supply chain (e.g., platforms, importers, distributors) can be held liable. However, some ambiguity remains around how and when AI systems are deemed “*defective*”.

The rCSAD

Key dates and facts:

- **Proposed:** 6 February 2024.
- **Entered into force:** Not yet; legislative procedure ongoing.
- **Transposition into Member State law by:** Three years after entry into force (see Art.33(1) rCSAD, as of 13 December 2024).
- **Last modified:** Proceedings in the Council, 13 December 2024.
- **Legal nature:** Binding EU Directive.

As part of the [EU Strategy for a more effective fight against child sexual abuse](#), the proposed rCSAD is aimed at enhancing and expanding the criminal law rules concerning child sexual abuse and exploitation, introducing stricter penalties for perpetrators and enhancing support for victims.

With regard to AI, the Commission's proposal presents several key amendments. First, the proposed revised definition of child sexual abuse material (CSAM) will include AI-generated content, including so-called "*deep fakes*". Second, the rCSAD increases the scope of criminal offences for solicitation of children for sexual purposes "*by means of information and communication technology*", which may include AI. Stakeholders have suggested a number of possible improvements, such as a clarification of the definition of CSAM with regard to AI-generated images.

Meanwhile, businesses that develop or use AI systems that are capable of creating CSAM (within the potentially very broad definitions noted above) will need to carefully consider their obligations under the EU AI Act.

The EDPS Orientations

Key dates and facts:

- **Published:** 3 June 2024.
- **Last modified:** Original publication by the EDPS on 3 June 2024.
- **Legal nature:** Non-binding EDPS guidance and observations.

The EDPS is effectively responsible for data protection by the EU institutions. The EDPS Orientations provide non-binding guidance to EU institutions on lawful personal data processing when using generative AI, aiming to help comply with the [Regulation \(EU\) 2018/1725](#) (the equivalent of the GDPR for EU institutions) which does not explicitly address AI. While the EDPS Orientations are not directly applicable to businesses, they contain helpful analysis of AI issues that arise in the context of EU data protection law.

The EDPS Orientations cover issues such as how to determine whether personal data is being processed during the various stages of a generative AI system's lifecycle; the importance and role of the data protection officer and of data protection impact assessments in the context of processing personal data using AI systems; compliance with principles such as data minimisation and data accuracy; legal bases for processing personal data using AI systems; and providing transparency to affected individuals. The EDPS Orientations also call for oversight, monitoring, and proper prevention mechanisms throughout the lifecycle of AI systems. These observations may help businesses to anticipate the approach that EU data protection authorities are likely to take in relation to AI.

Glossary

All article references are to the EU AI Act unless otherwise specified.

Advisory Forum	(Art.67) – The Advisory Forum is established to provide technical expertise and advice to the AI Office and the AI Board.
Affected person	(Recs.20 and 171; Art.2(1)(g)) – This term is not explicitly defined in the EU AI Act, but from context it appears to mean individuals affected by AI.
AI Board	(Arts.65 and 66) – The European Artificial Intelligence Board is a representative advisory board established to facilitate consistent and effective implementation of the EU AI Act.
AI literacy	(Rec.20; Art.3(56)) – Skills, knowledge, and understanding that allow providers, deployers, and affected persons – taking into account their respective rights and obligations in the context of the EU AI Act – to make an informed deployment of AI systems, as well as to gain awareness about the opportunities and risks of AI and possible harm it can cause.
AI Office	(Arts.3(47) and 64) – The Commission’s function of contributing to the implementation, monitoring, and supervision of AI systems and GPAI models, as well as AI governance. The EU AI Act states that references to the AI Office should therefore be construed as references to the Commission.
AI regulatory sandbox	(Art.3(55)) – A controlled framework set up by a competent authority which offers providers or prospective providers of AI systems the possibility to develop, train, validate, and test, where appropriate in real-world conditions, an innovative AI system, pursuant to a sandbox plan for a limited time under regulatory supervision.
AI system	<p>(Rec.12; Art.3(1)) – A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.</p> <p>See Chapter 3 (Core definitions) for further analysis of this term.</p>
AI systems guidelines	Guidelines published by the Commission on 6 February 2025 on determining whether a software system constitutes an AI system as defined in the EU AI Act.
Authorised representative	(Rec.82; Arts.2(1)(f) and 3(5)) – Any organisation located or established in the EEA who has received and accepted a written mandate from a provider of an AI system or a general-purpose AI model to, respectively, perform and carry out on its behalf the obligations and procedures established by the EU AI Act.

Biometric categorisation system	(Art.3(40)) – An AI system for the purpose of assigning natural persons to specific categories on the basis of their biometric data, unless it is ancillary to another commercial service and strictly necessary for objective technical reasons.
Biometric data	(Art.3(34)) – Personal data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of a natural person, such as facial images or dactyloscopic data.
Biometric identification	(Rec.15; Art.3(35)) – The automated recognition of physical, physiological, behavioural, or psychological human features for the purpose of establishing the identity of a natural person by comparing biometric data of that individual to biometric data of individuals stored in a database. See Chapter 3 (Core definitions) for further analysis of this term.
CE marking	(Art.3(24)) – A marking by which a provider indicates that an AI system is in conformity with the requirements set out in Arts.8 – 15, and other applicable EU harmonisation legislation providing for its affixing.
CJEU	The Court of Justice of the European Union.
Commission	The European Commission, being the main executive body of the EU.
Common specification	(Art.3(28)) – A set of technical specifications that prescribes technical requirements to be fulfilled by a product, process, service, or system, and which lays down one or more of requirements in Art.2(4)(a)-(d) of Regulation (EU) 1025/2012.
Conformity assessment	(Art.3(20)) – The process of demonstrating whether the requirements set out in Chapter III, Section 2, relating to a high-risk AI system have been fulfilled.
Conformity assessment body	(Art.3(21)) – A body that performs third-party conformity assessment activities, including testing, certification, and inspection.
Copyright in the Digital Single Market Directive	Directive (EU) 2019/790 on copyright and related rights in the Digital Single Market.
Council	Refers to the Council of Ministers of the EU, composed of government ministers from each EU Member State.
Database of high-risk AI systems	(Rec.131; Arts.71 and 49) – An EU database for high-risk AI systems to be set up and maintained by the Commission in collaboration with the Member States, containing certain information (see Annexes VIII and IX) on high-risk AI systems under Art.6(2), Annex III, and not-high-risk AI systems under Art.6(3), to be submitted by the respective provider or, in case of public authorities, the deployer.

Deep fake	(Art.3(60)) – AI-generated or manipulated image, audio, or video content that resembles existing persons, objects, places, entities, or events and would falsely appear to a person to be authentic or truthful.
Deployer	<p>(Rec.13; Art.3(4)) – A natural or legal person, public authority, agency, or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.</p> <p>See Chapter 3 (Core definitions) for further analysis of this term.</p>
Distributor	(Recs.83 and 84; Art.3(7)) – A natural or legal person in the supply chain, other than the provider or the importer, that makes an AI system available on the EEA market.
Downstream provider	(Art.3(68)) – A provider of an AI system, including a GPAI system, which integrates an AI model, regardless of whether the AI model is provided by the provider itself and vertically integrated, or provided by another entity based on contractual relations.
DMA	The Digital Markets Act. Regulation (EU) 2022/1925 on contestable and fair markets in the digital sector.
DSA	The Digital Services Act. Regulation (EU) 2022/2065, a regulation aimed at creating a safer digital space where the fundamental rights of users are protected and to establish a level playing field for businesses.
EDPS	(Art.70(9)) – The European Data Protection Supervisor is the data protection supervisory authority for the EU institutions, under Art.52 of Regulation (EU) 2018/1725. It is allocated additional supervisory powers in relation to AI, under the EU AI Act.
EEA	The European Economic Area, comprising EU Member States plus Iceland, Liechtenstein, and Norway.
Emotion recognition system	(Art.3(39)) – An AI system for the purpose of identifying or inferring emotions or intentions of natural persons on the basis of their biometric data.
e-Privacy Directive	Directive 2002/58/EC, an EU directive concerning the processing of personal data and the protection of privacy in the electronic communications sector.
FLOPs	The predominant methodology set out in the EU AI Act to evaluate GPAI model capabilities is the cumulative amount of computation used for training, measured in floating point operations (FLOPs) (see Chapter 12).

GDPR	General Data Protection Regulation (EU) 2016/679, a regulation on data protection and privacy for all individuals within the European Union and the EEA.
General Product Safety Regulation	Regulation (EU) 2023/988 on product safety.
GPAI model	<p>(Recs.97 – 99; Art.3(63)) – An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market, and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development, or prototyping activities before they are placed on the market.</p> <p>See Chapter 3 (Core definitions) for further analysis of this term.</p>
GPAI model with systemic risk	A GPAI model that meets the conditions of Art.51.
GPAI system	(Art.3(66)) – An AI system which is based on a GPAI model, and which has the capability to serve a variety of purposes, both for direct use as well as for integration in other AI systems.
GPAI code of practice	Pursuant to the GPAI code of practice, signatories commit to reproduce and extract only lawfully accessible copyright-protected content by not circumventing effective technological measures, excluding from their web-crawling piracy domains but also by complying with machine-readable protocols to express rights reservations.
Harm	(Rec.5) – Depending on the circumstances regarding its specific application, use, and level of technological development, AI may generate risks and cause harm to public interests and fundamental rights that are protected by Union law. Such harm might be material or immaterial, including physical, psychological, societal, or economic harm.
Harmonised standard	(Art.3(27)) – A standard adopted on the basis of a request made by the Commission for the application of EU harmonisation legislation (as defined in Art.2(1)(c) of Regulation (EU) 1025/2012).
Implementing acts	Non-legislative acts taken pursuant to specific rules contained in a legislative act. Implementing acts are normally taken by the Commission and are often of an administrative or technical nature.

Importer	(Recs.83 and 84; Art.3(6)) – Any organisation located or established in the EEA that places on the market an AI system that bears the name or trademark of a natural or legal person outside the EEA.
Intended purpose	(Art.3(12)) – The use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation.
In Vitro Diagnostics Regulation	Regulation (EU) 2017/746 on in vitro diagnostic medical devices.
Making available on the market	The supply of an AI system or a GPAI model for distribution or use on the EEA market in the course of a commercial activity, whether in return for payment or free of charge.
Market surveillance authority (MSA)	(Art.3(26)) – The national authority carrying out the activities and taking the measures pursuant to Regulation (EU) 2019/1020.
Market Surveillance Regulation	Regulation (EU) 2019/1020 on market surveillance.
MDR	Medical Devices Regulation (EU) 2017/745 on medical devices.
Member States	The 27 member states of the European Union, being Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, Poland, Portugal, Romania, Slovak Republic, Slovenia, Spain, and Sweden.
National competent authority	(Arts.3(48)) – A notifying authority or a market surveillance authority. As regards AI systems put into service or used by Union institutions, agencies, offices, and bodies, references to national competent authorities or market surveillance authorities are construed as references to the European Data Protection Supervisor.
NIS 2 Directive	The Network and Information Security Directive 2 (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union.
Notified body	(Art.3(22)) – A conformity assessment body notified in accordance with the EU AI Act or other relevant EU harmonisation legislation.

Notifying authority	(Art.3(19)) – The national authority responsible for setting up and carrying out the necessary procedures for the assessment, designation, and notification of conformity assessment bodies and for their monitoring.
Official Journal	The Official Journal of the European Union is the official publication (gazette) for EU legal acts, other acts, and official information from EU institutions, bodies, offices, and agencies.
Operator	(Rec.22; Art.3(8)) – A catch-all term for providers, product manufacturers, deployers, authorised representatives, importers, and distributors.
Placing on the market	(Art.3(9)) – The first making available of an AI system or a GPAI model on the EEA market.
Post-market monitoring system	(Art.3(25)) – All activities carried out by providers of AI systems to collect and review experience gained from the use of AI systems they place on the market or put into service for the purpose of identifying any need to immediately apply any necessary corrective or preventive actions.
Product manufacturer	(Rec.87; Art.2(1)(e)) – The concept of a “ <i>product manufacturer</i> ” is not explicitly defined in the EU AI Act (instead, it is defined in the EU harmonisation legislation listed in Annex I to the EU AI Act – see Rec.87). Product manufacturers are within the scope of the EU AI Act when they place an AI system on the EEA market together with their own products and under their own name or trademark.
Provider	<p>(Recs.21 and 22; Art.3(3)) – A natural or legal person, public authority, agency, or other body that develops an AI system or a GPAI model or that has an AI system or a GPAI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.</p> <p>See Chapter 3 (Core definitions) for further analysis of this term.</p>
Putting into service	(Art.3(11)) – The supply of an AI system, for first use directly to the deployer, or for the provider’s own use in the EEA, for its intended purpose.
Real-time remote biometric identification systems	(Art.3(42)) – A remote biometric identification system, whereby the capturing of biometric data, the comparison and the identification all occur without a significant delay, comprising not only instant identification, but also limited short delays in order to avoid circumvention.

Real-world testing plan	(Art.3(53)) – A document that describes the objectives, methodology, geographical, population, and temporal scope, monitoring, organisation, and conduct of testing in real-world conditions.
Remote biometric identification system	(Art.3(41)) – An AI system for the purpose of identifying natural persons, without their active involvement, typically at a distance through the comparison of a person’s biometric data with the biometric data contained in a reference database.
Risk	(Art.3(2)) – The combination of the probability of an occurrence of harm and the severity of that harm.
Safety component	(Art.3(14)) – A component of a product or of an AI system which fulfils a safety function for that product or AI system, or the failure or malfunctioning of which endangers the health and safety of persons or property.
Sandbox plan	(Art.3(54)) – A document agreed between the participating provider and the competent authority describing the objectives, conditions, timeframe, methodology, and requirements for the activities carried out within a sandbox.
SCD	Special category data – the categories of personal data listed in Arts.9 and 10 of the GDPR.
Scientific Panel	(Art.68) – The Scientific Panel of Independent Experts is to comprise members chosen by the Commission based on their scientific or technical expertise in the field of AI and their independence from any AI system providers.
Serious incident	(Art.3(49)) – An incident or malfunctioning of an AI system that directly or indirectly leads to any of the following: (a) the death of a person, or serious harm to a person’s health; (b) a serious and irreversible disruption of the management or operation of critical infrastructure; (c) the infringement of obligations under Union law intended to protect fundamental rights; (d) serious harm to property or the environment.
SMEs	Small and medium-sized enterprises.

Substantial modification	(Art.3(23)) – A change to an AI system after its placing on the market or putting into service which is not foreseen or planned in the initial conformity assessment carried out by the provider and as a result of which the compliance of the AI system with the requirements set out in Chapter III, Section 2, of the EU AI Act is affected or results in a modification to the intended purpose for which the AI system has been assessed.
Systemic risk	(Art.3(65)) – A risk that is specific to the high-impact capabilities of GPAI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain.
Testing in real-world conditions	(Art.3(57)) – The temporary testing of an AI system for its intended purpose in real-world conditions outside a laboratory or otherwise simulated environment, with a view to gathering reliable and robust data and to assessing and verifying the conformity of the AI system with the requirements of the EU AI Act. It does not qualify as placing the AI system on the market or putting it into service within the meaning of the EU AI Act, provided that all the conditions laid down in Arts.57 or 60 are fulfilled.
Transparency	(Rec.27) – The term is not explicitly defined in the EU AI Act; however, Rec.27 explains that the principle of transparency means that AI systems are developed and used in a way that allows appropriate traceability and explainability, while making humans aware that they communicate or interact with an AI system, as well as duly informing deployers of the capabilities and limitations of that AI system and affected persons about their rights.
VLOP	Very Large Online Platform, as outlined in Art.33 of the DSA.
VLOSE	Very Large Online Search Engine, as outlined in Art.33 of the DSA.

Contributors

Lead partners



Tim Hickman

Partner, London

E tim.hickman@whitecase.com



Sylvia Lorenz

Partner, Berlin

E sylvia.lorenz@whitecase.com



Jenna Rennie

Partner, London

E jenna.rennie@whitecase.com



Clara Hainsdorf

Partner, Paris

E chainsdorf@whitecase.com

Berlin



Sylvia Lorenz

Partner

E sylvia.lorenz@whitecase.com



Constantin Teetzmann

Local Partner

E constantin.teetzmann@whitecase.com



Katrin Helle

Counsel

E katrin.helle@whitecase.com



Galina Wedel

Associate

E galina.wedel@whitecase.com

Brussels



Geneva Forwood

Partner

E gforwood@whitecase.com



Assimakis Komninos

Partner

E akomninos@whitecase.com



Iakovos Sarmas

Associate

E iakovos.sarmas@whitecase.com



Jonas Huber

Associate

E jonas.huber@whitecase.com

London



Tim Hickman

Partner

E tim.hickman@whitecase.com



Rory Hishon

Partner

E rhishon@whitecase.com



Jenna Rennie

Partner

E jenna.rennie@whitecase.com



John Timmons

Partner

E john.timmons@whitecase.com

London *(continued)*



Di Yu
Partner

E di.yu@whitecase.com



Oli Bowley
Associate

E oli.bowley@whitecase.com



Elizabeth Hanson
Associate

E elizabeth.hanson@whitecase.com



Aishwarya Jha
Associate

E aishwarya.jha@whitecase.com



Sarah Lee
Associate

E sarah.lee@whitecase.com



Daniel Mair
Associate

E daniel.mair@whitecase.com



Jeffrey Shin
Associate

E jeffrey.shin@whitecase.com



Zarlush Zaidi
Associate

E zarlush.zaidi@whitecase.com

Los Angeles



Hope Anderson
Partner

E hope.anderson@whitecase.com

Madrid



Marcos Soberón
Local Partner

E marcos.soberon@whitecase.com

Paris



Clara Hainsdorf
Partner

E chainsdorf@whitecase.com



Alexandre Ghanty
Associate

E aghanty@whitecase.com

Prague



Anna Stárková
Associate

E anna.starkova@whitecase.com

Jon Gilbert (Professional Support Lawyer, White & Case, London), **Sulaiman Iqbal** (Trainee Solicitor, White & Case, Dubai), **Natasha Parsons** (Trainee Solicitor, White & Case, London), and **Parissa Santi-Weil** (Associate Director, White & Case, Paris) contributed to this publication.

About White & Case

White & Case is a global law firm with longstanding offices in the markets that matter today. Our on-the-ground experience, our cross-border integration and our depth of local, US and English-qualified lawyers help our clients work with confidence in any one market or across many.

We guide our clients through difficult issues, bringing our insight and judgment to each situation. Our innovative approaches create original solutions to our clients' most complex domestic and multijurisdictional deals and disputes.

By thinking on behalf of our clients every day, we anticipate what they want, provide what they need and build lasting relationships. We do what it takes to help our clients achieve their ambitions.

Band 1: Arbitration (International)

Chambers Global (Global Market Leaders) 2025

Band 1: Banking & Finance Band 1: Projects & Energy

Chambers Global (Global Multi-Jurisdictional) 2025

Structured Finance Deal of the Year

Airline Economics' Aviation 100 Global Leaders Awards 2025

Public M&A Team of the Year

British Legal Awards 2024

Catalyst Private Equity Deal of the Year

DealMakers Gala Awards 2024

Firm of the Year (Europe)

Chambers Europe Awards 2024

Digital Infra Financing of the Year

TMT M&A Awards USA 2024

#1 Infrastructure Legal Advisor for Transport (by deal value)

Infrastructure and Project Finance League Table 2024

Sovereign Debt Restructuring of the Year

Global Restructuring Review Awards 2024

Team of the Year – Project Finance

IFLR Asia-Pacific Awards 2024

Innovation in Generative AI Strategy

Financial Times Innovative Lawyers North America Award 2024

whitecase.com

In this publication, White & Case means the international legal practice comprising White & Case LLP, a New York State registered limited liability partnership, White & Case LLP, a limited liability partnership incorporated under English law, and all other affiliated partnerships, companies and entities.

This publication is prepared for the general information of our clients and other interested persons. It is not, and does not attempt to be, comprehensive in nature. Due to the general nature of its content, it should not be regarded as legal advice.

ATTORNEY ADVERTISING.

Prior results do not guarantee a similar outcome.

Portions of this publication were created with the assistance of AI tools, under human review and editorial control.

© 2025 White & Case LLP