MASON
HAYES &
CURRAN

JANUARY 2026

Artificial Intelligence

Annual Review 2025

Dublin London New York San Francisco

MHC.ie

MASON HAYES & CURRAN

Welcome Artificial Intelligence Annual Review 2025

Welcome to our Artificial Intelligence (AI) Annual Review for 2025. It was a momentous year for those following the developments in the AI space.

AI Act obligations came into effect

The first key compliance obligations were introduced in February, with the ban on prohibited practices under Article 5 and the AI literacy obligation under Article 4 taking effect. The obligations regarding general-purpose AI models came into effect in August, and it remains to be seen whether the next set of obligations for high-risk AI will kick in, in August 2026.

Guidelines and codes published

Key guidelines and codes have also been published this year, including:

- Prohibited AI practices guidelines
- Guidelines on the definition of AI systems
- AI literacy FAQs
- Guidelines on the scope of obligations for providers of general-purpose AI models
- The General-Purpose AI Code of Practice
- Draft serious incidents guidance

We expect some uncertainty in 2026 following the publication of the draft Digital Omnibus package. The package will have a significant impact on compliance planning given the uncertainty surrounding the high-risk AI deadline and the possibility for the AI literacy obligation to be removed for providers and deployers. With the enforcement powers of the AI Office coming into effect in August, we expect to see it begin to exercise its oversight and enforcement powers by issuing requests for information and potentially commencing more formal investigations.

In the meantime, we take this opportunity to look back at the key developments in the AI Act space that occupied our time in 2025.

In our AI Annual Review, we focus on:

- The guidelines on prohibited AI practices and its impact for AI providers
- The AI systems definition guidelines
- Key takeaways from the AI literacy FAQs
- The key obligations for those who chose to sign up to the General-Purpose AI Code of Practice
- What you need to know about the EU Commission's General-Purpose AI Model Guidelines
- The draft Serious AI Incidents Guidance
- Key learnings from the draft Digital Omnibus Package
- What we can expect in the year ahead

EDITORS

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

OISÍN TOBIN
Partner, Data & Technology
otobin@mhc.ie

HANNAH PERRY
Partner, Data & Technology
hperry@mhc.ie

Contents

Latest AI Insights from our Team

01

Commission Guidelines on Prohibited AI Practices Established

Clarifying scope, enforcement, and impact for AI providers

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

The Commission published its highly anticipated Guidelines on prohibited artificial intelligence practices on 4 February 2025.

The Guidelines are non-binding, but will help organisations determine whether their AI systems may be classified as prohibited AI under the AI Act.

It is crucial that organisations classify AI systems as falling in or out of scope of Article 5, given that non-compliance with this ban can result in significant fines for companies: 35 million euro or up to 7% of a company's annual worldwide turnover, whichever is higher.

WHAT YOU NEED TO KNOW

- Prohibited AI practices are already enforceable: Even though formal enforcement begins on 2 August 2026, the prohibitions under Article 5 have direct effect and can be enforced in court now, including via interim injunctions
- General-purpose AI systems can be caught: Providers must ensure their systems are not reasonably likely to behave or to be used for prohibited practices under Article 5, even if not built for a specific use
- Detailed scope and examples: The Guidelines offer clarifications, including what falls in and out of scope of the Article 5 prohibitions
- Review and prepare: Organisations should assess their AI systems against the Guidelines and update internal documentation to demonstrate compliance or out-of-scope classification

Overview

The Guidelines are comprehensive and cover:

- Relevant background and objectives of the Guidelines
- A general overview of prohibited AI practices including the scope of Article 5, enforcement, and the application of the prohibitions to general-purpose AI systems
- Article 5(1)(a) - Article 5(1)(h): Each section provides useful examples and considers for each *"article"*
  - The rationale and objectives
  - Main concepts and components
  - Scope, and
  - The interplay with other EU legislation
- Safeguards and conditions for the exceptions under Article 5(2) - 5(7)
- The entry into application of the prohibitions, and
- The status of the Guidelines

Key takeaways

Our five key takeaways are:

01 The Guidelines appear to be pragmatic on the whole and provide useful content on the scope of the provisions, exemptions and market definitions. They should be carefully reviewed by organisations in the context of their own AI systems, particularly regarding what might now fall out of scope.

02 Many of the prohibitions contain several cumulative conditions, all of which must be fulfilled for the prohibition to apply. Careful consideration should be given to each condition as there may be scope for the prohibition to be disapplied.

03 The Guidelines should be used by organisations in the context of preparing internal compliance documentation which can be used to demonstrate why their AI systems are out of scope of Article 5.

04 The AI Act should not be reviewed in isolation. For each prohibition, guidance is provided on the interplay with other EU laws. When reviewing AI systems for potential prohibited practices, due consideration should be given to other EU laws that may impact the application and scope of the prohibition.

05 The Guidelines provide helpful context in terms of the rationale and objectives of the prohibitions. It appears clear that these prohibitions are not intended to operate as strict liability provisions. Instead, they need to be carefully considered in view of the context, purpose and objective of the prohibition.

Can general-purpose AI systems be prohibited AI systems?

While there has been some ambiguity as to whether a general-purpose AI system can be subject to Article 5, the Guidelines confirm that the ban on prohibited AI practices can be applicable to these types of AI systems, ie those which do not have a specific intended purpose. The Guidelines advise that providers have a responsibility not to place on the market/put into service AI systems, including general-purpose AI systems, that are "*reasonably likely to behave or be directly used*" in a manner that would be in contravention of Article 5.

The Guidelines point to a number of measures and built-in safeguards that providers are expected to put in place to prevent and mitigate harmful behaviour and misuse by deployers. These measures should be implemented provided they are feasible and proportionate, having regard to the AI system itself and the context.

The key challenge for providers of general-purpose AI systems is how to avoid being in scope of Article 5 where the prohibition under Article 5 is linked to a very specific purpose of the system, ie Articles 5(1)(d)-(h) of the AI Act.

The Guidelines helpfully acknowledge that in these cases providers of general-purpose AI systems have limited options to avoid these prohibited practices and advise that providers will have to rely primarily on:

- Excluding prohibited use of AI systems in contracts with deployers, ie terms of use
- Provide appropriate information in the instructions of use for deployers and regarding the necessary human oversight

The Guidelines also note that in certain circumstances monitoring for compliance with that restriction on prohibited uses may be appropriate. It does not, however, clarify in what circumstances it might be appropriate to implement this monitoring.

Other measures providers may implement, mentioned in the Guidelines, include:

- Appropriate, safe and ethical design
- Integration of technical and other safeguards
- Restrictions of use
- Transparency and user control, and
- Appropriate information in the instructions of use

The measures suggested in the Guidelines lack specificity and necessary detail. However, it provides some indication of what measures the AI Office might expect providers of general-purpose AI systems to take to avoid placing an AI system on the market that is or could become in scope of Article 5.

Scope of prohibited AI practices

The Guidelines set out in detail the scope of the various prohibited practices and provide illustrative examples. We have highlighted some key points and included some helpful examples from the Guidelines below:

1. Article 5(1)(a) and (b) – harmful manipulation, deception and exploitation

The Guidelines deal with Article 5(1)(a) and 5(1)(b) together and acknowledge that there is interplay between both articles. It clarifies that the key difference between the two is that Article 5(1)(a) is primarily focused on the nature of the techniques and the covert nature of the AI system's influence which can undermine an individual's ability to make free choices. Article 5(1)(b) is mostly focused on the protection of vulnerable individuals who may be more susceptible to exploitation by AI systems.

The Guidelines provide detailed guidance on how key terms under Articles 5(1)(a) should be interpreted, which include the following concepts:

- Subliminal techniques: These are described as techniques that "*operate beyond (below or above) the threshold of conscious awareness*". An example of this is a visual subliminal message, ie a message/image that flashes briefly during video playback which is technically visible, but appears too quickly for the conscious mind to register, which can influence individuals' attitude/behaviour.
- Purposefully manipulative techniques: Techniques that are designed or aim to *"influence, alter, or control an individual's behaviour*" in a way that undermines their *"individual autonomy"* and "*free choices*". An example of this is sensory manipulation, eg background audio that leads to mood alterations.
- Deceptive techniques: This should be understood in accordance with recital 29 of the AI Act. An example of this is a chatbot that pretends to be the friend of a person using a synthetic voice and scams that person causing significant harm.

It also provides further details on core concepts under Article 5(1)(b) such as the concept of materially distorting behaviour, and the meaning of *"harm"* under both articles.

Out of scope examples

Recital 29 of the AI Act acknowledges that *"common and legitimate commercial practices"* such as advertising, should not be regarded in and of itself as harmful manipulative AI enabled practices.
The Guidelines helpfully provide further information on how to distinguish between lawful commercial practices which are based on persuasion, and those practices which are prohibited. It is crucial to understand the difference between these two concepts given that lawful persuasion is a permissible practice and outside the scope of the prohibition.

The Guidelines recognise that while both manipulation and persuasion influence decisions and behaviours, they differ significantly. The key differentiators are:

- Transparency
- Objective and impact of the technique
- Consent, and
- Compliance with legal and regulatory frameworks

In addition, the Guidelines provide some examples of AI systems that are out of scope of Articles 5(1)(a) and (b) on the basis that they do not cause significant harm:

- An AI companionship system: *"designed in an anthropomorphic way and with affective computing to make the system more appealing and effectively makes users more engaged, but is not engaging in other manipulative or deceptive practices in a manner that is reasonably likely to cause them serious psychological, physical or other harms, unhealthy attachment and dependency."*
- A therapeutic chatbot: which *"uses subliminal techniques to steer users towards a healthier lifestyle and to quit bad habits, such as smoking. Even if the users who follow the chatbot's advice and subliminal therapy experience some physical discomfort and psychological stress due to the effort made to quit smoking, the AI-enabled chatbot cannot be considered likely to cause significant harm. Such temporary discomfort is unavoidable and outweighed by the long-term benefits for users' health. There are no hidden attempts to influence decision-making beyond promoting healthy habits."*

2. Article 5(1)(c) – social scoring

The Guidelines suggest that certain legitimate, beneficial and justified purposes for the evaluation or classification of persons, including those to improve the effectiveness of processes, quality of service, safety etc may mean that the AI system falls outside the scope of this prohibition. This is particularly the case where there is compliance with applicable legislation and appropriate safeguards implemented.

Out of scope examples

The Guidelines provide examples of legitimate scoring practices in line with EU and national law that are outside the scope of the prohibition:

- Fraud detection: *"Companies have a legitimate interest to evaluate customers for financial fraud and those practices are not affected by the prohibition, if the evaluation is based on relevant data such as transactional behaviour and metadata in the context of the services, past history and other factors from sources that are objectively relevant to determine the risk of fraud and if the detrimental treatment is justified and proportionate as a consequence of the fraudulent behaviour".*
- Profiling for safety: *"Online platforms profiling users for safety reasons on their services based on data which is relevant for the context and purpose of assessment is out of scope of Article 5(1)(c) AI Act, when the evaluation does not result in detrimental treatment that is disproportionate to the gravity of the user's misbehaviour".*

3. Article 5(1)(d) – individual risk assessment and prediction of criminal offences

This prohibited practice concerns AI systems used to make risk assessments of individuals to assess or predict their likelihood of committing a criminal offence. In order to be in scope of this prohibition, it must be that based *"solely"* on profiling or assessments based on personality traits and characteristics. An assessment may not be based solely on profiling if the AI system is used to support a human assessment. Providers should note that if the AI system is used to support a human assessment then it is not prohibited, but will be deemed a high-risk AI system under Annex III, point 6(d).

While the Guidelines do not rule out the prohibition applying to private actors, as opposed to just law enforcement authorities, it suggests that private actors may be caught where they are entrusted by law enforcement to exercise public authority and public powers.

Out of scope examples

Some out of scope examples mentioned in the Guidelines include:

- Location-based or geospatial predictive or place-based crime predictions: *"A customs authority uses AI risk analytic tools to predict the likelihood of the location of narcotics or illicit goods, for example on the basis of known trafficking routes."*
- AI systems that support human assessments based on objective and verifiable facts linked to a criminal activity: *"The use of an AI system that assesses the risk whether a prisoner should receive the benefit of an early release. The AI profile of the affected person or assessment of their personality traits and characteristics only support the human assessment of objective and verifiable facts related to past criminal offences and demonstrated behaviour relevant to rehabilitation."*

4. Article 5(1)(e) – untargeted scraping of facial images

Article 5(1)(e) prohibits the use of AI systems that create or expand facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage.

The Guidelines provide further clarity on key concepts under this prohibition, including the following terms:

- Facial recognition databases: These databases are capable of matching a human face, from images/video, against a database of faces, and compares the two to identify whether there is likely a match. It may be temporary, centralised or decentralised. The Guidelines state that this prohibition applies to any database that can be used for facial recognition, regardless of whether this is the "*sole purpose*" of the database or not.
- Untargeted scraping: The Guidelines note that the wording of Article 5(1)(e) implies that this prohibition does not apply to any scraping tool, but that it specifically applies only to tools for untargeted scraping. Untargeted scraping is the indiscriminate scraping of data, ie not aimed at a particular individual/group of individuals.

Out of scope examples

The Guidelines provide useful examples of instances where the prohibition would not be applicable, including:

- The untargeted scraping of biometric data other than facial images, such as voice samples
- Where no AI systems are involved in the scraping of facial images
- Where facial image databases are not used for the recognition of persons

5. Article 5(1)(f) – emotion recognition

The Guidelines confirm that the emotion recognition prohibition is only applicable where biometric data is used to infer emotions. Further guidance is provided on what might constitute biometric data. For example, this may include:

- Physiological biometrics such as fingerprints, contours of their face, and
- Behavioural biometrics such as walking, keystrokes, eye tracking and heartbeats

The prohibition applies to the workplace and educational institutions, both public and private and potentially vocational schools. A key feature is that education institutions provide a certificate, and participation is a precondition for obtaining the certificate. An AI based app using emotion recognition for learning a language online outside an education institution is not prohibited. However, if an education institution mandates the use of the app, then it is prohibited.

Out of scope examples

- AI systems that infer emotions/sentiments not on the basis of biometric data
- AI systems that infer *"physical states"* such as *"pain and fatigue"*
- Emotion recognition systems used in all other domains other than in the areas of the workplace and education institutions

6. Article 5(1)(g) – biometric categorisation for certain 'sensitive' characteristics

In order for this prohibition to apply, the AI system must have the objective of deducing or inferring a limited number of sensitive characteristics including: race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation.

The Guidelines confirm that to fall outside the scope of the definition of biometric categorisation, two conditions must be met:

- It is ancillary to another commercial service, and
- Strictly necessary for objective technical reasons

Out of scope examples

The Guidelines note that the prohibition does not cover AI systems involved in the labelling or filtering of lawfully acquired biometric datasets. It provides the following examples of labelling/filtering that would be out of scope:

- *"The labelling of biometric data to avoid cases where a member of an ethnic group has a lower chance of being invited to a job interview because the algorithm was 'trained' based on data where that particular group performs worse, ie has worse outcomes than other groups."*
- *"The categorisation of patients using images according to their skin or eye colour may be important for medical diagnosis, for example cancer diagnoses."*

7. Article 5(1)(h) – real-time remote biometric identification (RBI) systems for law enforcement purposes

The Guidelines reiterate the core components of Article 5(1)(h) and provide some further guidance on key definitions related to this prohibition:

- Remote biometric identification systems: The definition implies the lack of active involvement of the subject. In other words, where there is no active participation. It results in capturing the person's characteristics, typically at a distance. The identification aspect involves the comparison of biometric data that has been captured with biometric data contained within a database, eg a criminal database.
- Real-time: This means that data is captured near instantaneously or without *"significant delay*".
- Publicly accessible spaces: Does not include online spaces, eg social media.
- For the purpose of law enforcement: Law enforcement purposes are given a wide definition, which include investigation, detection, prosecution and crime prevention.

As noted in the Guidelines, there are certain exceptions to this prohibition where this practice is strictly necessary for the objectives set out in Article 5(1)(h)(i)-(iii) of the AI Act.

02

Defining AI

Commission guidelines on AI systems

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

LEONA CHOW
Associate, Data & Technology
lchow@mhc.ie

The EU Commission published their much anticipated guidelines on the definition of an artificial intelligence system on 6 February 2025. The guidelines explain how the legally defined term *"artificial intelligence system"* is applied in practice.

In particular, the guidelines aim to assist providers in determining whether a software system constitutes an AI system.

In this article, we provide an overview of the guidelines.

WHAT YOU NEED TO KNOW

- The EU Commission published non-binding guidelines on how to interpret the definition of an AI system under the AI Act, on 6 February 2025
- The definition is broken into seven elements, including autonomy, inference, and the ability to influence environments, highlighting inference as an important aspect
- Techniques such as machine learning, and logic and knowledge based approaches are in scope, while techniques such as basic data processing and simple prediction systems are out of scope
- The guidelines recommend first classifying an AI system in accordance with its risk category under the AI Act to determine if it is out of scope, before considering whether it meets the definition of an AI system

Scope of application

The guidelines specifically state that they are designed as a guide only and do not provide an exhaustive list of all AI systems that may be covered. They are not legally binding, and any authoritative interpretation of the AI Act can ultimately only be provided by the Court of Justice of the European Union.

Breaking out the definition

Essentially, the guidelines break down the definition into its seven main elements and provide detailed explanations for each. The seven elements are that the system is:

1. *a machine-based system;*
2. *that is designed to operate with varying levels of autonomy;*
3. *that may exhibit adaptiveness after deployment;*
4. *and that, for explicit or implicit objectives;*
5. *infers, from the input it receives, how to generate outputs;*
6. *such as predictions, content, recommendations, or decisions;*
7. *that can influence physical or virtual environments.*

Pre and post-deployment included

Importantly, the guidelines note that the definition adopts a lifecycle-based perspective encompassing two main phases:

(i) The pre-deployment or 'building' phase of the system, and

(ii) The post-deployment or 'use' phase of the system, referencing a recent OECD paper[1] on the same topic

This approach is highlighted to clarify that the seven elements of the definition are not required to be present continuously throughout both phases of that lifecycle. Instead, the definition acknowledges that specific elements may appear at one phase, but may not persist across both phases. This is an important point for those looking to make precise scoping arguments. It reflects a means of analysis deployed in recent data protection supervisory authority guidelines.

In-scope and out-of-scope

Prior to the guidelines' publication, most commentators focused on two or three crucial aspects of the definition that go to the heart of what does and does not constitute an AI system. The most important aspects were seen as autonomy and inference, with many also including adaptiveness.

Reading between the lines, it seems the Commission has zoned in on inference as the key aspect of the definition. Almost six of the thirteen pages of the guidelines are devoted to this topic and the majority of the guidelines focus on listing the AI techniques that fall within the scope of the definition. It also outlines techniques that may fall outside the scope, such as comparing AI software with simple execution or rules-based software.

In-scope techniques are:

1. Machine learning approaches including:
   - Supervised learning
   - Unsupervised learning
   - Self-supervised learning
   - Reinforcement learning
   - Deep learning

2. Logic and knowledge based approaches including:
   - Knowledge representation
   - Inductive (logic) programming knowledge bases
   - Inference and deductive engines
   - Symbolic reasoning
   - Expert systems, and
   - Search and optimisation methods

Out-of-scope techniques are:

- Systems for improving mathematical optimisation, including linear or logistic regression methods
- Basic data processing
- Systems based on classical heuristics, and
- Simple prediction systems

How to use these guidelines

In the final section, the guidelines explain how they should be used when determining whether a system is considered an AI system under the AI Act. According to the guidelines, this assessment should be based on the specific design and function of the system taking into account the seven key elements of the definition.

In our view, the guidelines are most helpful to those with AI systems that are founded on a technique specifically identified as out-of-scope, or those who have a very specific query on scope. Organisations looking to make a quick big picture call on "*in v out of scope*" of the AI Act are not best served by beginning with assessing their technology against these guidelines, given how broadly the guidelines interpret the AI systems definition.

As recommended in the guidelines, the optimal approach for assessing whether your organisation may be subject to the AI Act is to take the following steps. First, consider how the use of the technology might be classified under the AI Act, such as whether it could fall into a high-risk category. It may be the case that there will be no compliance lift, for example, if it is a minimal risk AI system. If it is likely to fall under one of the higher risk categories such as high-risk AI, the second step is to consider whether the system is excluded from the scope of the AI Act altogether on the basis that it does not meet the definition of an AI system in the first place.

[1] OECD (2024), "Explanatory memorandum on the updated OECD definition of an AI system", OECD Artificial Intelligence Papers, No. 8, OECD Publishing, Paris, https://doi.org/10.1787/623da898-en, p.7.

Enforcement and entry into application

The Guidelines provide useful clarifications on enforcement more generally under the AI Act.

It clarifies that, where there are cross border implications beyond a market surveillance authority's (MSA) territory, it must inform the Commission and the other MSAs. It notes that all Member States must follow the union safeguard procedure under Article 81, with a decision taken by the Commission. This approach aims to ensure uniformity of the prohibitions across the EU, and to provide legal certainty for providers and deployers. The Commission also notes that MSAs should "*strive for a harmonized application*" of the prohibitions for comparable cases in other Member State territories by drawing inspiration from the Guidelines and cooperating within the AI Board.

The Guidelines further note that the *ne bis in idem* principle contained within Recital 168 should be respected. This principle concerns the provision of multiple penalties for the "*same prohibited conduct*". It helpfully provides the example of non-labelling deep fakes, which may also constitute a deceptive technique under Article 5(1)(a).

Enforcement before August 2025

The Guidelines also acknowledge that the provisions on enforcement and penalties will not apply before 2 August 2025, nor will the MSAs be set up before then. However, it notes the prohibitions are applicable in the interim period before the provisions on enforcement/penalties apply. It further notes that these prohibitions have *"direct effect"* and therefore enable affected individuals/parties to enforce the prohibitions in court and by requesting interim injunctions.

Next steps

We recommend that organisations take the time to review the Guidelines, which are highly detailed and provide useful clarification and out of scope examples.

Organisations should carefully consider whether any further review of their AI systems is warranted, including whether any internal documentation or procedures need to be updated.

> *The prohibitions have 'direct effect' and therefore enable affected individuals/parties to enforce the prohibitions in court and by requesting interim injunctions.*

03

AI Literacy Takeaways

Key takeaways from the AI literacy FAQs

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

Article 4 of the AI Act requires that AI providers and deployers must ensure a sufficient level of AI literacy among their staff. This means ensuring that individuals have the level of skill, knowledge and understanding that allows them to:

- Make an informed use of AI systems
- Have awareness about the risks and possible harm AI can cause, and
- Have an awareness of the opportunities of AI

Article 4 applies to providers and deployers of AI systems and general-purpose AI systems. However, it does not apply to general-purpose AI models (GPAI Models). Instead, Article 4 is addressed to any organisation that uses any AI system. AI literacy should be ensured for all staff, and *"any other person"* dealing with AI systems, ie *"affected persons"*.

Guidance

The Commission published its 'AI Literacy - Questions & Answers' (the FAQs) in May 2025 to clarify what is expected of organisations providing or deploying AI systems. We set out a summary of the key takeaways from the FAQs below.

Key requirements

The FAQs state that strict requirements will not be imposed by the AI Office. Rather, organisations will be afforded a degree of flexibility in determining what constitutes a *"sufficient level"* of literacy. However, as a minimum, AI literacy training should:

- Ensure a general understanding of AI across the organisation. This means staff should understand:
  - What is AI?
  - How does it work?
  - What AI is used in the organisation?
  - What are its opportunities and dangers?
- Take into account the organisation's role, e.g. is it a provider or deployer of AI systems? This is because an organisation's role under the AI Act will influence the required literacy level.
- Have regard to the risk of the AI system that is being provided or deployed. Organisations must assess the educational needs of their staff based on the use of AI systems. Organisations must also provide training on the risks associated with its use and any mitigations users should be aware of.
- Consider the following points as part of the implementation of the programme :
  - Differences in technological knowledge, experience, education, and training amongst those taking part in the literacy programme, and
  - The context in which the AI systems will be used and the individuals they will impact e.g. what sector the AI system will be used in, and its purpose.

Enforcement

Article 4 of the AI Act entered into force on 2 February 2025. As a result, the obligation to ensure a sufficient level of AI literacy already applies. The relevant supervision and enforcement powers will take effect on 2 August 2026. Enforcement will be the responsibility of the National Market Surveillance Authorities, or *"MSAs"*.

Public enforcement therefore cannot commence until those MSAs have been designated and granted the necessary powers, which many Member States have so far have been delayed in doing. However, the FAQs also note the potential for *"private enforcement"* by persons who suffer harm due to an organisation's failure to comply with its literacy obligations. While the affected individual could sue according to national law, the FAQs point out that the AI Act does not provide for criminal offences or indeed, a right to compensation.

Penalties

While the AI Act does not provide specific penalties for breach of the AI literacy obligation under Article 99, the FAQs clarify that MSAs could impose penalties and other enforcement measures to sanction infringements of Article 4. Any enforcement would be based on national law, which only a small number of Member States have implemented so far. The FAQs note however that any sanction must be proportionate and take into account factors such as the nature and gravity of the infringement, as well as its intentional or negligent character. It suggests an infringement might be more likely if there is proof of an incident due to a failure to provide appropriate training to employees.

Possible reforms

The Commission, as part of its Digital Omnibus, is set to propose a raft of reforms to the AI Act. These include a potential reversal of the obligations related to AI literacy. If adopted, the proposal would replace Article 4 entirely. Under the new Article 4, the Commission and Member States would be responsible for encouraging AI literacy, rather than enforcing an obligation to do so against providers and deployers. According to the leaked draft, this particular reform has been proposed with the explicit intention of reducing the compliance burden imposed by the AI Act.

Next Steps

Organisations may take comfort that there is no one-size-fits-all approach to AI literacy. While reports of reform are to be welcomed, there is still significant uncertainty as to whether the European Parliament will embrace the Commission's simplification agenda. We would therefore still recommend reviewing the use of AI within your organisation as well as any training, policies or procedures governing its use. Irrespective of whether the AI literacy obligation for providers/deployers is removed, we recommend carrying out some form of AI literacy training in any event, given that it is good practice and helps to mitigate any risks associated with AI use.

> *Organisations will be afforded a degree of flexibility in determining what constitutes a 'sufficient level' of literacy... but training should, at a minimum, ensure a general understanding of what AI is and how it works.*

04

General-Purpose AI Code of Practice

Setting the standard for general-purpose AI model providers

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

LEONA CHOW
Associate, Data & Technology
lchow@mhc.ie

The European Commission published the final draft of the General-Purpose AI Code of Practice (the Code) in July 2025. The Code is voluntary and consists of three chapters. Chapter I (Transparency) and II (Copyright) apply to all signatories. Chapter III (Safety and Security) is only relevant to providers of models with systemic risk.

While not legally binding, providers of GPAI Models can rely on the Code to demonstrate compliance with their obligations under Articles 53 and 55 of the AI Act. Providers who do not sign up will be required to demonstrate compliance by alternative adequate means.

WHAT YOU NEED TO KNOW

- The Commission's final draft of the General-Purpose AI Code of Practice offers a voluntary route for general-purpose AI model, or *"GPAI Model"* providers to demonstrate compliance with their obligations under the AI Act.
- Signatories must meet detailed transparency requirements, including maintaining a single documentation form, updating it after relevant changes, cooperating with downstream providers, and retaining records for 10 years.
- The Code sets clear copyright standards for lawful data gathering, compliance with text-and-data-mining opt-outs, mitigation of infringing outputs and providing a complaints mechanism for rightsholders.
- Providers of GPAI Models with systemic risk must meet additional safety and security commitments, including developing a *"state-of-the-art"* framework, assessing and mitigating risks throughout the lifecycle, and reporting serious incidents to the AI Office.

Transparency

The transparency chapter describes the measures that signatories of the Code commit to implementing as part of their transparency obligations under Articles 53(1)(a) and (b). The Code provides a *"Model Documentation Form"* which enables signatories to include all of the information required in a single place.

Providers should bear in mind that in order to comply with the transparency chapter, signatories must ensure that all measures are adhered to, including:

- **Updates/retention:** the Model Documentation Form must be updated to reflect any *"relevant changes"*. Previous versions of the Model Documentation Form must also be retained for a period of 10 years after their initial creation.
- **Contact information:** contact information should be publicly disclosed on the provider's website or other appropriate means for the AI Office, or downstream providers, to request access to the relevant information in the Model Documentation Form.
- **Cooperating with downstream providers:** downstream providers should be furnished with the most up-to-date documentation intended for downstream providers. They should also

be provided with additional information where requested where this information is necessary for the downstream provider to have a good understanding of the capabilities and limitations of the GPAI Model relevant to its integration into AI systems. The information should be provided within a reasonable timeframe and no later than 14 days after the request was made, apart from *"exceptional circumstances"*.

- **Quality control:** the documented information must be controlled for quality and integrity, retained as evidence of compliance with the AI Act. It should also be protected from unintended alterations.

Copyright

While the AI Act provides little guidance on what exactly should be contained within the copyright policy, the copyright chapter of the Code provides detailed measures that signatories must implement to demonstrate compliance with Article 53(1)(c).

Copyright policy

Signatories must adhere to the following requirements in the context of the copyright policy:

- **Single document:** the copyright policy must be in a single document, which incorporates the measures set out in the copyright chapter.
- **Oversight:** signatories are required to assign responsibilities within their organisation for implementing and overseeing the policy.

Using only lawfully accessible copyright-protected content when crawling

The copyright chapter requires signatories that are making use of web-crawlers to scrape data ensure that:

- Technological measures intended to prevent or restrict unauthorised access to protected works, as defined in Article 6(3) of Directive 2001/29/EC, are not circumvented, particularly by way of pay-walls or subscription restrictions, and
- Websites known to persistently and repeatedly publish content that infringes on intellectual property rights are excluded from their web-crawling activities. The copyright chapter states that a dynamic list of relevant websites will be compiled to assist with this.

Complying with rights reservations when crawling

The copyright chapter also provides that signatories must comply with text and data mining opt-outs. In particular, signatories must:

- Only use web crawlers that read and follow instructions expressed in accordance with Robot Exclusion Protocol, or *"robots.txt"*, or any subsequent version of this
- Comply with other appropriate machine-readable protocols to express opt-outs

Mitigating copyright infringing outputs

The risk of copyright infringement via outputs must be mitigated by:

- The implementation of appropriate and proportionate technical safeguards preventing the GPAI Model from reproducing protected works
- Prohibiting copyright infringing uses of the model, e.g. via an acceptable use policy, terms and conditions, etc.

Lodging complaints

Finally a mechanism must be put in place, via which rightsholders can submit complaints concerning the signatory's non-compliance with their commitments under the Code. A point of contact for electronic communication with those rightsholders must also be appointed.

Safety and security

The final chapter relates the obligations of providers of GPAI Models with systemic risk under Article 55. It is the lengthiest chapter with a total of 10 commitments:

- **Commitment 1:** requires that providers of GPAI models with systemic risk must develop a *"state-of-the-art"* "Safety and Security Framework" (the Framework). The Framework must outline the processes and measures which will be put in place to assess and mitigate the systemic risk attached to the model. Providers must ensure that their Framework is in place no later than four weeks after notifying the Commission that their model meets the systemic risk threshold, and no later than two weeks before placing the model on the market. Signatories should reassess the Framework where the provider has reasonable grounds to believe that it is no longer adequate. Commitment 1 states that reassessment should be done:
  - Following serious incidents or near misses indicating that unacceptable risks have occurred, or
  - Every 12 months after the model is placed on the market
- **Commitment 2:** provides various measures regarding how signatories should identify systemic risk.
- **Commitment 3:** outlines how signatories should carry out systemic risk analysis.
- **Commitment 4:** requires signatories to set out systemic risk acceptance criteria and to determine whether the systemic risks stemming from the model are acceptable.
- **Commitment 5:** requires the signatories to implement appropriate safety mitigations along the entire model lifecycle to ensure the systemic risks stemming from the model are acceptable.
- **Commitment 6:** requires signatories to implement an adequate level of cybersecurity protection for their models and their physical infrastructure along the entire model lifecycle, in accordance with the measures listed in the Code.
- **Commitment 7:** requires signatories to report to the AI Office information about the model, systemic risk assessment and mitigations by creating a *"Safety and Security Model Report"*.
- **Commitment 8:** requires signatories to define clear responsibilities for managing systemic risks across all levels of the organisation.
- **Commitment 9:** outlines the appropriate processes and measures signatories must comply with to keep track of, document and report serious incidents to the AI Office.
- **Commitment 10:** outlines the additional transparency measures and documentation signatories are required to draw up and keep up to date. These include detailed descriptions of the model's architecture, its integration into AI systems, model evaluations, and safety mitigations implemented. Signatories are also required to publish a summarised version of their Safety and Security Framework and Model Report(s), as well as any other updates to the extent *"necessary"* to assess or mitigate systemic risks.

> *Providers of GPAI models with systemic risk must develop a 'state-of-the-art' Safety and Security Framework outlining the processes and measures put in place to assess and mitigate risks throughout the entire model lifecycle.*

05

EU Commission General-Purpose AI Model Guidelines

Navigating the new guidance

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

LEONA CHOW
Associate, Data & Technology
lchow@mhc.ie

The Commission adopted formal guidelines in July 2025 on the scope of the obligations for general-purpose AI models under the AI Act, *"the Guidelines"*. The finalised version largely reflects the draft guidelines discussed in our 2025 Mid-year Review, though a number of important changes were made during the drafting process.

WHAT YOU NEED TO KNOW

- The European Commission has issued final guidelines clarifying how the AI Act applies to general-purpose AI models, or *"GPAI Models"*, including when a model qualifies as a GPAI Model and who is responsible for compliance.
- A new *"lifecycle"* approach means GPAI Model providers must keep documentation, copyright policies and training data summaries updated throughout a model's development and use.
- Providers of GPAI Models with systemic risk may be expected to notify the AI Office before model training is complete.
- Third parties who modify a GPAI Model may become GPAI Model providers if they reach the training compute threshold of one-third of the training compute of the original GPAI Model.
- The guidelines also explain the rules for placing models on the EU market, conditions for the open-source exemption, transitional arrangements and how the AI Office will supervise and enforce compliance.

Overview

The Guidelines cover the following topics:

- What is a GPAI Model, including what is a new distinct model versus a modified version
- Who is the provider of a GPAI Model, including when a downstream provider becomes subject to GPAI Model provider obligations
- What constitutes placing on the market and the criteria for the open-source exemptions
- Methods for estimating the computational resources used to train or modify a model
- Transitional rules, grandfathering, retroactive compliance, and supervision and enforcement of the GPAI Model obligations

The Guidelines on GPAI Models are non-binding. This is because authoritative interpretation may only be given by the Court of Justice of the European Union. Despite their non-binding nature, they provide important clarifications on how the AI Office will interpret and apply the obligations under the AI Act. In particular, the AI Office notes its exclusive responsibility for the supervision and enforcement of the obligations of providers of GPAI Models.

The Guidelines are expected to evolve over time and will be updated as necessary, particularly in light of evolving technological development.

What is a GPAI Model

The Guidelines do not provide a specific list of tasks or capabilities to help providers classify their models. Rather, the AI Office's approach to assess whether a model qualifies as a GPAI Model is based on the computational resources used to train it, as well as the modalities of the model. According to the AI Office an *"indicative criterion"* of whether a model is a GPAI Model is where:

- The training compute exceeds 1023 FLOP, and
- The model can generate language as text or audio, or it is text-to-image or text-to-video

According to the AI Office, this threshold approximately corresponds to the *"amount of compute"*, or computing power, typically used to train a model with one billion parameters on a large amount of data. Helpfully, the AI Office acknowledges that even where models meet the requirements of training compute and modality set out above, where it *"exceptionally"* does not display significant generality or cannot competently perform a wide range of distinct tasks, then it is not a GPAI Model. On the other hand, where a model does not meet the indicative criterion, it may still be a GPAI Model where it displays significant generality and is capable of performing a wide range of distinct tasks. For example, the AI Office notes that an AI model that:

- has a training compute above the 1023 FLOPs threshold; and
- can generate text

is not a GPAI Model, despite meeting the indicative criterion, because it is only capable of transcribing speech, and not a wide range of distinct tasks.

GPAI Models with systemic risk

The Guidelines discuss when models may be classified as a GPAI Model with systemic risk. In accordance with Article 51(1) of the AI Act, a model will be classified as having systemic risk where it has *"high-impact capabilities"*. Essentially, these are capabilities that match or exceed those recorded in the most advanced models. Alternatively the Commission may designate a model as having high-impact capabilities in accordance with the criteria set down in Annex XIII. Once a model is classified as having systemic risk, the provider must notify the AI Office *"without delay and in any event within two weeks after that requirement is met or it becomes known that it will be met"*.

Notably, the Guidelines state that this notification may be required *before* training is complete as long as the provider can reasonably foresee that their model is likely to qualify as having high-impact capabilities. According to the AI Office, given that the planning and *"upfront allocation of compute resources"* will take place before any training-run occurs, providers should estimate the cumulative amount of training compute they will use. They will then need to notify the Commission where that estimate meets the threshold set down in Article 51(2), currently 10^{25} FLOP. In addition, providers are expected to closely monitor both the actual and expected compute usage over the course of a model's training and throughout the model's lifecycle.

The lifecycle of a GPAI Model

The Guidelines introduce the concept of the *"lifecycle"* of a model. According to the AI Office, the lifecycle of a GPAI Model begins at the start of the large pre-training run, and it is to be interpreted in a broad sense. Any subsequent development of the model, by the same provider or on its behalf, forms part of the same model lifecycle and does not give rise to new or distinct models.

In practice this means:

- The documentation that must be drawn up under Article 53(1)(a) and (b) must be kept up to date throughout the model's entire lifecycle.
- The copyright policy under Article 53(1)(c), must be applied throughout the entire lifecycle of the model. Providers may also develop one policy and apply it to all models.
- The summary of content used for training under Article 53(1)(d) must also be updated in accordance with the training data template provided by the AI Office during the lifecycle of the AI model. The explanatory note to the training data template, published on 24 July 2025, says it must be updated whenever the provider further trains the GPAI Model on new data that requires changes to the summary. This update must be performed either every six months or sooner if the new data is *"materially significant"*.
- For GPAI Models with systemic risk, the systemic risk assessment and mitigation under Article 55(1) should be carried out continuously throughout the model's entire lifecycle.

Downstream modifiers

The Guidelines set out when third parties who modify a GPAI Model, known as *"downstream modifiers"*, will themselves be treated as GPAI Model providers. The AI Office considers that a downstream modifier will only become the provider of the modified GPAI Model where it leads to a *"significant change"* in the model's generality, capabilities, or systemic risk.

It notes that an indicative criterion is where the training compute used to modify the model is greater than a third of the original model's training compute.

In situations where the down-stream modifier cannot be expected to know this value, the Guidelines permit the modifier to replace that threshold with:

- One-third of the training compute for a model presumed to be a GPAI Model, currently 10^{23} FLOP
- One-third of the training compute for a model presumed to have *"high-impact capabilities"*,currently 10^{25} FLOP, if the original model is a GPAI Model with systemic risk

The AI Office confirms that the downstream modifier is only responsible for the GPAI Model provider obligations regarding the modified aspect. In particular, the documentation required under Article 53(1)(a) and (b) is limited to information concerning the modification. In addition, the copyright policy under Article 53(1)(c) and summary of training data under Article 53(1)(d) only needs to take into account data used as part of the modification.

If a downstream modifier modifies a GPAI Model with systemic risk, the resulting model will also be presumed to have high-impact capabilities.

The downstream modifier will be required to comply with the obligations applicable to providers of GPAI Models with systemic risk, and must notify the Commission in accordance with Article 52(1).

Placing on the market

The Guidelines provide various examples of placing GPAI Models on the market. Examples include:

- Making a GPAI Model available via an API
- Copying a GPAI Model onto a customer's own infrastructure
- Integrating a GPAI Model into a chatbot and making it available via a web interface
- Uploading a GPAI Model to a public catalogue, hub, or repository for direct download on the EU market
- Use for internal processes that are essential for providing a product or service to third parties, or that affect the rights of natural persons in the EU

The AI Office notes that these examples should be interpreted in accordance with the Blue Guide.

Notably, the Guidelines note that where an upstream actor makes a model available outside the EU market to a downstream actor, and that downstream actor integrates the model into an AI system and places that AI system on the EU market, the model shall be considered to be placed on the market. Where the upstream actor does not exclude the model's supply/distribution on the EU market, including via integration into AI systems, in a *"clear unequivocal way"* then the upstream actor will be considered to be the provider of that model under the AI Act and will be subject to the relevant obligations. Conversely, where the downstream actor breaches these clear and unequivocal terms by placing that AI system (which integrates the model) on the market, then they will be considered the provider of the model.

Open-source models

The Guidelines elaborate on the requirements to meet, in order to benefit from the open-source exemption under Article 53(2), namely:

- Conditions on the licence
- Lack of monetisation, and
- Public availability of parameters, including the weights, the information on the model architecture, and the information on model usage

> *Providers of GPAI Models with systemic risk may be expected to notify the AI Office before model training is complete.*

Conditions on the licence

The Guidelines explain that *"free and open-source"* means that the license should allow wide dissemination of the model and incentivise further development. The licence must allow users to freely access, use, modify and redistribute the model. According to the AI Office, the following restrictions would disqualify the licence from meeting this criteria:

- Limitations to non-commercial or research-only use
- Prohibitions on distributing the model or its components
- Usage restrictions triggered by user scale thresholds, and
- Requirements to obtain separate commercial licences for specific use cases

Lack of monetisation

In order for the exemption to be applicable, no monetary compensation should be required in exchange for access, use, modification, and distribution of the GPAI Model. The Guidelines provide various examples of *"monetisation"*, including:

- Users being required to purchase support, training, and maintenance services to access the model
- The model being provided under a dual licensing model or similar approach that allows for example, free academic use, but requires payment for commercial usage or use over a certain scale, and
- Technical support or other services that are indistinguishably linked to the model itself, which require payment, and without which the model would not work or be accessible

Public availability of information

According to the AI Office information about the parameters, including the weights, model architecture, and model usage that should be available publicly in accordance with Article 53(2), should have a degree of clarity. It should also be sufficiently specific to allow access, usage, modification, and distribution of the AI model. The information about the model's usage should, according to the Guidelines, contain at least the following information:

- Information about the model's input and output modalities
- Capabilities and limitations
- The technical means, e.g. instructions for use, infrastructure, tools, required for the model to be integrated into AI systems, which may include the appropriate configuration for the intended use cases, where applicable

This is necessary to ensure that the model can be used by other parties for practical applications.

Grandfathering provision

The Guidelines confirm that GPAI Models that benefit from the grandfathering provisions do not require re-training or unlearning. This applies where:

- Retraining or unlearning is not possible for actions performed in the past
- Where information on training data is not available, or
- Retrieval would be disproportionate for the provider

In these cases, this must be clearly justified and disclosed in the copyright policy and summary of the content used for training.

The AI Office suggests that model providers who are placing a GPAI Model on the market after 2 August 2025, especially GPAI Models with systemic risk, and foresee difficulties with complying with their obligations, should proactively inform the AI Office how, and when, they will take the necessary steps to comply with their obligations. Importantly, the Guidelines suggest that providers of GPAI Models with systemic risk are still required to notify the AI Office, even though the AI Office cannot take any enforcement actions until August 2026.

General-Purpose AI Code of Practice

The Guidelines state that signatories of the General-Purpose AI Code of Practice *"the Code"* will *"benefit from increased trust from the Commission and other stakeholders."* The Guidelines confirm that, for signatories, the AI Office is expected to focus its enforcement activities on monitoring whether signatories are adhering with the Code. For non-signatories, the Guidelines note that these providers will be expected to report the measures they have implemented to the AI Office. They will also need to explain how they comply with their obligations under the AI Act via other adequate means, such as by carrying out a gap analysis.

The Guidelines seem to suggest that signatories of the Code will be subject to more favourable treatment. For instance, the Guidelines state that non-signatories may be subject to more requests for information and requests for access to conduct model evaluations. In addition, it appears to suggest that signatories may be subject to reduced fines, given that the Commission may take into account commitments implemented in line with the Code as a mitigating factor when fixing the amount of fines.

Enforcement

The Guidelines confirm that the AI Office will supervise and enforce the obligations for GPAI Model providers and AI systems based on those GPAI Models where they have the same providers. In addition the Guidelines clarify that:

- The AI Office will take a *"collaborative and proportionate approach"*. It encourages close informal cooperation with all providers during the training phase of the GPAI Model to facilitate compliance and to ensure market placement without delays, and in particular for those GPAI Models with systemic risk.
- The AI Office expects proactive reporting by providers of GPAI Models with systemic risk, whether providers are signatories of the Code or not.
- The obligation to report serious incidents under Article 55(1)(c) covers *"serious cybersecurity breaches related to the model or its physical infrastructure, including the (self-)exfiltration of model parameters and cyberattacks"*

Next steps

We recommend that all GPAI Model providers and potential downstream modifiers should carefully consider these Guidelines. While these Guidelines provide helpful clarifications they may in some instances go beyond what is necessary or required under the AI Act.

> *The AI Office encourages close informal cooperation with providers during the training phase to facilitate compliance.*

06

Serious Incidents

Commission publishes guidance on reporting obligations

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

ALISON STENSON
Senior Associate, Data & Technology
astenson@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

Under the AI Act, providers of high-risk AI systems will be required to report serious incidents to national authorities. This obligation, set out in Article 73, aims to detect risks early, ensure accountability, enable quick action and build public trust in AI technologies. In anticipation of implementing the rules, the Commission produced draft guidance and a reporting template for review and feedback. This was done by way of public consultation which ended in early November 2025.

It is important for those in the high-risk AI space to be aware of these guidelines as these rules and reporting templates are expected to apply as early as Q3 2026. As noted, it would have been preferable to have the high-risk AI guidelines in place before this stage.

WHAT YOU NEED TO KNOW

- The Commission's draft guidelines confirm that Article 73 applies only to high-risk AI systems. General-purpose AI model (GPAI Model) related incidents are addressed in the General-Purpose AI Model Code of Practice and GPAI Model guidelines.
- A *"serious incident"* covers a wide range of events, including malfunctions, misclassifications, major drops in accuracy, downtime or unexpected behaviour. They must directly or indirectly cause, or be likely to cause, significant harm.
- Providers must report incidents to the relevant national authority, investigate them, carry out a risk assessment and take corrective action. Deployers also have reporting duties and must immediately notify providers and market authorities when a serious incident is identified.

The draft guidelines are very much in draft form with no fewer than five placeholders, taking the form of *"For example: […]"* for incomplete sections. This indication of haste is not something we have seen to date with other AI Act guidelines.

Just high-risk AI?

The draft guidelines apply to high-risk AI systems only and confirm this is the scope of Article 73. The confirmation here is helpful as there was some speculation based on loose drafting that it could apply to more than high-risk AI systems.

The draft guidelines also confirm that they do not apply to GPAI Models and specifically Article 55(1)(c) which deals with serious incidents in relation to systemic risk models. For serious incidents concerning GPAI Models under the relevant Article, the Code of Practice facilitates demonstrating compliance by way of its Commitment 9 in the Safety and Security Chapter. The guidelines on the scope of the obligations for providers of GPAI Models also provide further context regarding the scope of the serious incidents definition in the context of GPAI Models.

The scope is not limited to high-risk AI system serious incidents, as defined in Article 3(49). It also covers widespread infringements of high-risk AI systems, as defined in Article 3(61).

Scope - a signal of pragmatism

Those working in the chatbots and GenAI spaces can take some comfort from the confirmation that serious incident reporting applies only to providers and deployers of high-risk AI systems. Section 2 of the guidelines sets out the constituent parts of what we should understand a *serious incident* to be, especially as incident is not defined in the Act. The guidelines seem to draw inspiration from existing sectoral EU regulatory frameworks and an OECD definition for scoping for this:

"An incident is a not planned/programmed deviation in the characteristics of performance. OECD defines an AI incident as an event where the development or use of an AI system results in actual harm, while an event where the development or use of an AI system is potentially harmful is termed an "AI hazard".

The definition of a serious incident also includes a malfunctioning of the AI system giving us the following combined list of practical examples for incidents / malfunctions:

- Misclassifications
- Significant drops in accuracy
- Temporary system downtime
- Unexpected system behaviour

A causal link

The incident or malfunction needs to be causal, or likely to be causal and it is, if without it, the harm in its concrete form would not have occurred. The causation can also be indirect, i.e. secondary effects. Indirect examples include:

- *An AI system provides an incorrect analysis of medical imaging, leading a physician to make an incorrect diagnosis or treatment decision, which subsequently causes harm to the patient.*
- *An AI based credit scoring system incorrectly flags the unreliability of a person and a loan is denied based on this decision.*

As a noteworthy aside, many of the examples provided relate to medical device type high-risk AI, HR and credit scoring.

Breaking down the definition

Other aspects of the ingredients of the definition of serious incident are also reviewed:

1. *Directly or indirectly caused*
2. *Death of a person or serious harm to a person's health*
3. *Serious and irreversible disruption of the management or operation of critical infrastructure*
4. *Infringements of obligations under Union law intended to protect fundamental rights,*
5. *"Serious" harm to property, and*
6. *"Serious" harm to the environment*

Each is worth a look, especially for those who will operate in specific verticals where poor outcomes from uses of their high-risk AI systems will lead to risks in the certain spaces e.g. HR AI systems in the fundamental rights space.

Who is obliged to report

Article 73 (1) AI Act specifies that providers of high-risk AI Systems need to report any serious incidents to the market surveillance authorities of the Member States where those incidents occurred. If the exact location is not known to the provider, it is the business location of the deployer that counts.

Providers are also obliged to perform investigations regarding the serious incident and the AI system concerned. Any investigation must include a risk assessment of the incident, and corrective action, as provided by Article 73 (6) of the AI Act. In addition, providers must cooperate with the competent authorities, and where relevant with notified bodies concerned during the investigations.

Deployers have a reporting role also. Where they have identified a serious incident, they must immediately inform the provider, and then the importer or distributor as well as the relevant market authorities as set out under Article 26 (5).

What's next?

Organisations should keep a close eye on the publication of the Commission's final guidelines, which is expected in Q3 2026. Although the current draft is not yet final, it offers a basic outline of the Commission's current thinking. The draft guidelines can be used proactively to benchmark existing processes, identify capability gaps, and to integrate serious incident reporting into existing incident-response and escalation procedures, for example under established cybersecurity, GDPR, and product-safety incident frameworks, which can significantly streamline future compliance workloads.

As the regulatory landscape continues to evolve, organisations that begin preparing now will be better positioned to respond efficiently to serious incidents, mitigate potential harm, and meet expectations from regulators.

07

Digital Omnibus Package

Key Learnings

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

The long-awaited draft digital omnibus package was published in mid-November 2026. It aims to streamline and simplify the AI Act. For organisations providing or deploying AI in the EU market, understanding the draft digital omnibus package is an essential part of ensuring AI Act readiness.

It offers an early opportunity to map obligations and any changes, anticipate how supervisory authorities will coordinate, and identify where simplification may provide practical compliance benefits. We outline the key elements of the draft digital omnibus package as they relate to the AI Act. We also examine what these developments mean for providers and deployers operating in this space.

The focus is on HRAI not GPAI Models

As expected, the focus of the proposals will be on high-risk AI (HRAI) rather than general-purpose AI models *"GPAI Models"*. There are a few interesting proposals on the GPAI Model side to be cognisant of.

1. AI Office's role

Indeed, there is only one proposal that applies directly to the regulation of GPAI Models, and it's a governance change to Article 56 on the GPAI Codes of Practice. Article 56(6) of the AI Act sees the AI Office and Board jointly monitor and evaluate the effectiveness of the Codes of Practice. In the proposal, the control of this process passes to the Commission in its own right with only an obligation to take on board the opinion of the Board.

> *"The Commission and the Board shall regularly monitor and evaluate the achievement of the objectives of the codes of practice by the participants and their contribution to the proper application of this Regulation. The Commission, taking utmost account of the opinion of the Board, shall assess whether the codes of practice cover the obligations provided for in Articles 53 and 55, and shall regularly monitor and evaluate the achievement of their objectives. The Commission shall publish its assessment of the adequacy of the codes of practice."*

2. Changes to Article 75 (Enforcement)

Article 75 now explicitly confirms that the AI Office shall *"be exclusively competent for the supervision and enforcement of"* AI systems built on GPAI Models where the provider of each is the same. The same now applies to AI systems that *"constitute or that are integrated into a designated very large online platform or very large search engine"* within the meaning of the DSA. This *"exclusivity"* will be welcome news for the GenAI industry, which may now potentially avoid enforcement and requests for information from multiple market surveillance authorities *"MSAs"*.

In addition, the centralised enforcement powers of the AI Office are now extended to *"a large number of AI systems built on general-purpose AI models or embedded in very large online platforms and very large search engines"*.

3. Changes to Article 77

Similar to Article 75, the Commission wants to make a change here to clarify that the Article 77 bodies can only access provider information from the MSA and not the provider. In the process, however, they have dropped the references to high-risk AI. This means that this right could now apply to information related to GPAI Models, especially when one considers that the AI Office acts as an MSA under Article 75. It also seems to open the door to a broader concept of sharing of information between these bodies.

HRAI delay

Providers will welcome the Commission's desire to see a delay to the obligations related to HRAI systems. The Commission states, rather vaguely, that it will be *"linking the implementation timeline of high-risk rules to the availability of standards or other support tools"*. Based on the highly-publicised delays on the publication of the standards, it suggests that the HRAI obligations will indeed be delayed beyond the August 2026 deadline, possibly even later.

The revised draft of Article 113 seems designed so that any postponement cannot extend beyond 2 December 2027, although it may be set for an earlier date. The Commission wants the delay to be tied to the date of its decision to adopt the standards for HRAI that are currently in the drafting process.

- If that decision date falls before 2 August 2026, the delay will be pushed only by 6 months from the date of that decision for Annex III HRAI, i.e. the end of 2026 at the earliest.
- If the Commission adopts the standards after 2 August 2026, the deadline moves out but cannot extend beyond 6 months from July 2027. The current projected date for the standards to be ready for review by the Commission is end 2026 / beginning 2027. Factoring in a few months for a review and adoption phase, together with the six-month period set out in the Omnibus proposal, places the likely date in Q3 2027.

The Commission's complicated proposal on this issue will reassure its critics. It argues that if the standards are delivered quickly, which is unlikely, though this is not the Commission's fault, the HRAI deadline would only need to be delayed by the minimum amount of time.

The proposed delay to Article 50(2)

It is proposed that the application of fines for both the AI systems (Article 99) and GPAI Models (Article 101) sections concerning the Article 50(2) transparency marking obligations be delayed for one year, until 2 August 2027. However, MSAs will still be able to enforce those transparency obligations from 2 August 2026 but only for those systems placed on the EU market after 2 August 2026. Those placed on the market prior to 2 August 2026 have a grace period of six months from enforcement of the Article 50(2) obligations also. So, for systems launched after 2 August 2026, it is conceivable that the ultimate sanction of an order to recall an AI system or GPAI model from the market can be imposed during the period 2 August 2026 to 2 August 2027, even though fines cannot yet be applied.

Given the guidelines on these transparency provisions are not expected to be published until well into 2026, this is still a very tight timeline for those in the GenAI industry in particular.

Placing on the market clarified?

Those organisations engaging with our AI team on the concept of *"placing on the market"*, and how this applies to HRAI already on the market before 2 August 2026, will be familiar with the concern that the grace period for these product lines remains vulnerable. The issue is that the Blue Guide is clear on what *"placing on the market"* means. It applies to every product, not the product line as a whole. In other words, each time an AI system is made available to a new customer, it could be considered newly placed on the market. It is not treated as a one-off event linked to the product's first launch. That creates issues if you are relying on launching a HRAI prior to 2 August 2026 to extend the deadline for compliance, i.e. until there is a significant design change. The proposal aims to address this by saying that there needs to be a derogation from that Blue Guide interpretation. In this context, the concept of *"placing on the market"* should mean a broader category that covers a type and model of a HRAI system, and not each individual unit of that HRAI system.. It remains to be seen whether this interpretation of *"placing on the market"* is applied more broadly in the AI Act to AI systems in general.

Interaction with DSA

The explanatory notes of the proposal state that there needs to be some changes to the AI Act to simplify how the AI Office and the Commission should work together in relation to AI systems embedded in or qualifying as a very large online platforms or search engines. It says that the first point of entry for the assessment of the AI systems is the risk assessment, mitigating measures and audit obligations prescribed by Articles 34, 35 and 37 of the DSA. This is without prejudice to the AI Office's powers to investigate and enforce non-compliance with the rules of the AI Act. It seems the AI Office, under the AI Act, and Commission, under the DSA, will be expected to frequently cooperate regarding their activities relating to AI systems they both regulate. This approach is aimed at ensuring cooperation and avoiding overlapping fines. However, this arrangement of sharing information, as currently envisaged, seems to be informal . It proposes that the respective authorities can only use the information for the purposes of supervision, and enforcement can only take place if the undertaking agrees.

The Article 6(3) derogation

The punishing obligation for a provider to register their decision to rely on an Article 6(3) derogation is to be removed. The proposal notes that it constitutes a disproportionate compliance burden. However, it might instead impose an obligation on the provider to document the assessment before that system is placed on the market, or put into service. This assessment may be requested by national competent authorities. This lines up with our advice to put in place justification documents to ensure you are ready to deal with MSAs and the AI Office, especially when an edge case is identified.

A new look Article 4 and SCD processing

The Article 4 obligation regarding AI literacy looks set to change. Instead of imposing a direct duty on providers and deployers, it may be removed and replaced with an obligation for the Commission and Member States to rather *"encourage"* providers and deployers to ensure their staff have an adequate level of AI literacy. In time, this could result in Member States issuing soft guidelines on how to implement literacy training, which could result in an AI literacy lite model.

The Article 10(5) special category data processing derogation is being moved to Article 4 and extended to all AI systems and models. This makes sense and will be welcomed by GPAI model providers in particular.

Future guidance to be published confirmed

The updated draft also provides a list of guidance that will be published in the future, which providers will welcome. This includes:

- Guidelines on the practical application of the high-risk classification
- Guidelines on the practical application of the transparency requirements under Article 50 AI Act
- Guidance on the reporting of serious incidents by providers of high-risk AI systems
- Guidelines on the practical application of the high-risk requirements
- Guidelines on the practical application of the obligations for providers and deployers of high-risk AI systems
- Guidelines with a template for the fundamental rights impact assessment
- Guidelines on the practical application of rules for responsibilities along the AI value chain
- Guidelines on the practical application of the provisions related to substantial modification.
- Guidelines on the post-market monitoring of high-risk AI systems.
- Guidelines on the elements of the quality management system which SMEs and SMCs may comply with in a simplified manner.
- Guidelines on the AI Act's interplay with other Union legislation. By way of example:
  - Joint guidelines of the Commission and European Data Protection Board on the interplay of the AI Act and EU data protection law
  - Guidelines on the interplay between the AI Act and the Cyber Resilience Act, and
  - Guidelines on the interplay between the AI Act and the Machinery Regulation

Guidelines on the competences and designation procedure for conformity assessment bodies to be designated under the AI Act.

What's next?

While it is important to consider the draft omnibus package carefully and take note of any changes, it is worth noting that it still needs to go to the Parliament and Council for approval. After that step, there will be some Trilogue engagement before it can ultimately be agreed. This process will take a number of months at a minimum. It is possible that the Parliament will push back on any attempts to *"water down"* the existing law. So, bear in mind that the draft omnibus package is at best a starting position, not the final outcome. In any event, it is important to track this draft and follow any developments closely given that these changes may have a significant impact on any planned AI Act compliance strategies, in particular for high-risk AI.

08

Looking ahead to 2026

From AI Act readiness to full compliance

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

The regulatory landscape for AI will reach a critical inflection point in 2026, but also one marked by uncertainty. The majority of the obligations under the AI Act are scheduled to apply from 2 August 2026, critically with enforcement powers both at national and EU level kicking in for both models and systems. However, with the introduction of the draft Digital Omnibus Package by the European Commission, enforcement of high-risk AI system obligations remains uncertain.

Key developments on the horizon

- **Main compliance deadline 2 August 2026:** From that date, most of the AI Act's substantive requirements including transparency and high-risk AI obligations are in effect, subject of course to the proposed changes under the Omnibus Package.

- **Enforcement begins:** Market Surveillance Authorities and the AI Office will be empowered to supervise compliance and impose sanctions for non-compliance. All eyes will be on the first enforcement of the general-purpose AI model regime and what approach the AI Office will take. Will it send requests for information and documentation to all model providers, or just those making available systemic risk models? Or perhaps will enforcement arise as a result of a downstream provider complaint? There are a number of ways enforcement could arise.

- **Uncertainty amid regulatory changes:** The introduction of the Digital Omnibus Package, currently being advanced by the European Commission, may alter key aspects of the AI Act's implementation and impact compliance planning for providers and deployers of AI models and systems. For example, under the Omnibus proposals:

  - The deadlines for high-risk AI systems under Annex III may be postponed, to as late as December 2027 This will depend on when technical requirements, standards and conformity-assessment frameworks are finalised.
  - The AI literacy obligations for providers and deployers under Article 4 could be removed or significantly softened, with the burden placed instead on the Commission or Member States to encourage AI literacy.
  - There may be an introduction of a grace period of six months for the transparency obligations under Article 50(2) for those that have placed their AI systems on the market prior to 2 August 2026.

It's not all negative though. Those in the GenAI space will welcome the proposed changes that confirm the AI Office as the exclusive regulator of not only general-purpose AI models but also AI systems built on them by the same provider.

These proposals remain under discussion at EU level, so critically the shape of the AI Act regime as of August 2026 may look different from what many organisations have been planning for.

What this means for providers and deployers

- **Plan for compliance but stay agile:** companies should prepare as though the August 2026 deadline will be firm, including preparing for high-risk AI compliance, transparency measures and governance structures etc. However, companies should also track developments under the Digital Omnibus Package, since shifting deadlines or relaxed obligations may materially alter risk assessments.

- **Don't assume a delay or grace period will materialise for high-risk AI:** a grace period is by no means guaranteed as negotiations are still underway and transitional delays may not apply across all high-risk systems.

- **Enforcement risk is real in 2026:** for many companies, 2026 will be the first moment when failure to comply could result in real legal consequences, from regulatory penalties to reputational and operational risk. It is essential to ensure compliance is in place.

Strategic considerations for 2026

Use the year ahead as a *"last window"* for readiness

Companies providing or deploying AI systems, especially high-risk AI, should view early 2026 as the crucial window to complete inventory, map risk and prepare for compliance. Those in the model space need to be ready for enforcement and how to deal with AI Office requests for information and documentation.

Monitor legislative developments closely

Given the potential for change under the Digital Omnibus Package, compliance strategies should remain flexible; companies should be prepared to pivot if obligations are delayed, amended or simplified.

Invest in governance and training now

Even where AI literacy obligations may be softened, a robust internal compliance culture will remain critical, not only for legal compliance, but also for operational integrity and reputational resilience.

In short, 2026 will be the year when AI regulation moves from "preparing for" to "having to comply with". The window for preparation is now and the cost of underestimating this regime may be high.

at the end -add back the dese
.select= 1
.select=1
.scene.objects.active = modifier_
"selected" + str(modifier_ob)) # modi
_ob.select = 0
.context.selected_objects[0]
.objects[one.name].select = 1
("please select exactly two objects
OPERATOR CLASSES
10 1 0
0

> *Compliance in 2026 isn't just about meeting a deadline; it's about staying agile enough to pivot as the Digital Omnibus Package reshapes the regulatory finish line.*

MASON HAYES & CURRAN

ABOUT US

Mason Hayes & Curran is a business law firm with 124 partners and offices in Dublin, London, New York and San Francisco.

Our legal services are grounded in deep expertise and informed by practical experience. We tailor our advice to our clients' business and strategic objectives, giving them clear recommendations.

This allows clients to make good, informed decisions and to anticipate and successfully navigate even the most complex matters.

Our service is award-winning and innovative. This approach is how we make a valuable and practical contribution to each client's objectives.

KEY CONTACTS

BRIAN MCELLIGOTT
Partner, Head of Artificial Intelligence
brianmcelligott@mhc.ie

OISÍN TOBIN
Partner, Data & Technology
otobin@mhc.ie

HANNAH PERRY
Partner, Data & Technology
hperry@mhc.ie

ALISON STENSON
Senior Associate, Data & Technology
astenson@mhc.ie

SADHBH MURPHY
Associate, Data & Technology
sadhbhmurphy@mhc.ie

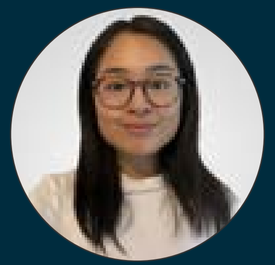

LEONA CHOW
Associate, Data & Technology
lchow@mhc.ie

For more information and expert advice, visit:

MHC.ie/AI