SLAUGHTER AND MAY

EU AI Act

Digital
HORIZON SCANNING

August 2026

Introduction

The EU's AI Act came into force on 1 August 2024, and has generally applied since 2 August 2026. Cited by the Commission as the 'world's first comprehensive AI law', the Act caught the attention of organisations across the globe with its focus on responsible innovation, wide extra-territorial reach, high fines and prohibition on certain types of AI which pose an unacceptable risk. It has had a phased transition, with most of the Act now applying.

There were also concerns from some within the EU that its cross-cutting provisions could negatively impact the EU's competitiveness and ability to foster innovation, and added unnecessary regulatory burdens on organisations. This led to its review and amendment as part of the EU's Omnibus (which itself was part of a wider simplification program targeting the EU's digital rules).

The AI act was reviewed and amended as part of the AI Omnibus. Its changes include delays to the application of the high risk and (some) transparency provisions, narrowing what is in scope for Annex 1 high-risk AI systems, removing duplication with sector rules, giving more authority to the AI Office and helping SMEs. See our blog for more details.

In this briefing, we highlight some of the key aspects of the EU AI Act ('the Act'), focussing on ten questions. We will also look at what practical steps organisations can take now to comply with the Act, and how the Act compares with the approach to AI regulation being taken in other jurisdictions like the UK and US.

Content

Note: while the Act is relatively detailed, more information on how it will work in practice is being provided through secondary legislation, harmonised standards, guidance and guidelines. Only some of these are currently available. For providers of AI systems and GPAI models, more details on the Act can also be found in our **article**

1. What is the aim of the Act?

The Act lays down a uniform legal framework for the development, marketing and use of artificial intelligence in line with EU values. It is a Regulation, meaning it has direct effect in EU member states.

The Commission states that the Act aims to foster responsible innovation in Europe – "[by] guaranteeing the safety and fundamental rights of people and businesses it will support the development, deployment and take-up of trustworthy AI in Europe."

As well as laying down rules around the development and use of AI, the Act contains provisions designed to promote innovation, including sandboxes and rules to ease the regulatory burden for SMEs.

2. What does the Act cover?

(i) AI Systems

While EU legislation in the technology space has tended to be technology neutral, the Act is specifically designed to regulate "AI Systems". However, defining AI Systems has proved challenging, with the original Commission proposal being widely criticised for being too broad. The final definition, which aligns with the approach proposed by the OECD, defines an AI System as:

> "a machine-based system 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."

Examples of AI Systems include: machine learning models such as autonomous driving systems; natural language processing such as chatbots, voice assistants and machine translation; computer vision systems such as facial recognition and medical imaging; and robotic process automation systems such as data entry systems and customer service bots.

Guidelines on the definition were published by the Commission in July 2025.

There are also specific rules for general-purpose AI ("GPAI") models, which are defined separately from AI Systems as AI models that display significant generality, are capable of competently performing a wide range of distinct tasks and that can be integrated into a variety of downstream systems or applications. For more information on GPAI models, see question 7 below.

(ii) Exemptions

While the Act is far reaching, some areas are out of its scope. It does not, for example, apply to areas outside the scope of EU law (e.g. it should not affect member states' competences in national security) and there are a number of exemptions listed, in areas such as military/defence, R&D and third country public authority use. It does not apply to any research, testing (other than in real world conditions) or development activity which takes place prior to the AI System being placed on the market or put into service, and users are also not caught if they are using AI purely for personal use.

The Act also provides an exemption for AI Systems released under free and open-source licences, unless those AI Systems trigger the criteria for prohibited or high-risk AI, or AI requiring additional transparency obligations.

Finally, the majority of the Act does not apply to operators of high-risk AI Systems (for more information on high-risk AI Systems see question 5 below) that have been put on the market or into service before those rules apply unless significant changes are made to their design (see questions 6 and 10 for more detail).

3. Who is caught by the Act's obligations?

The Act covers a wide range of organisations in the AI supply chain, including:

- Providers – organisations that develop and/or supply an AI System or GPAI model, or that have an AI System or GPAI model developed, and place it on the market or put it into service under their own name or trademark (whether for payment or free of charge).
- Deployers – organisations that use an AI System, where that use is under the deployer's authority (except if for personal/non-professional use).
- Other members of the supply chain, including Product Manufacturers (who incorporate an AI System into their product design and thereby place an AI System on the market or into service with their product, under their own name or trademark), Importers and Distributors.

Most of the Act's obligations apply to either Providers, who must effectively design and provide safe AI, or Deployers, who must use AI in a responsible and safe way.

4. Does the Act have extra-territorial reach?

Yes.

First, the Act covers Providers – wherever based – who place AI Systems or GPAI models on the market in the EU or put them into service in the EU. Non-EU organisations would therefore be subject to the Act if (among other things) they sell tools or services using AI in the EU.

Providers of high-risk AI Systems who are subject to the Act but not based in the EU must appoint an Authorised Representative who is located in the EU. The recitals state that this Authorised Representative:

"plays a pivotal role in ensuring the compliance of the high-risk AI Systems placed on the market or put into service in the Union by those providers who are not established in the Union and in serving as their contact person established in the Union."

Second, the Act covers non-EU Providers and Deployers where the outputs produced by their AI Systems are used in the EU. This may be difficult to determine in practice and organisations will therefore need to consider how they will monitor where the outputs from their systems are being used.

5. How does the Act's risk-based approach work?

The Act takes a risk-based approach to regulation. It defines four categories of risk and imposes separate obligations on Providers of GPAI models (as defined in question 2 above; see also question 7 for more information).

The obligations an organisation will face differ depending on both the role that organisation plays in the AI supply chain and the type of AI System involved.

The risk categories are set out below.

(i) Prohibited AI practices:

These have an unacceptable level of risk and include things like social scoring or an AI System which: deploys

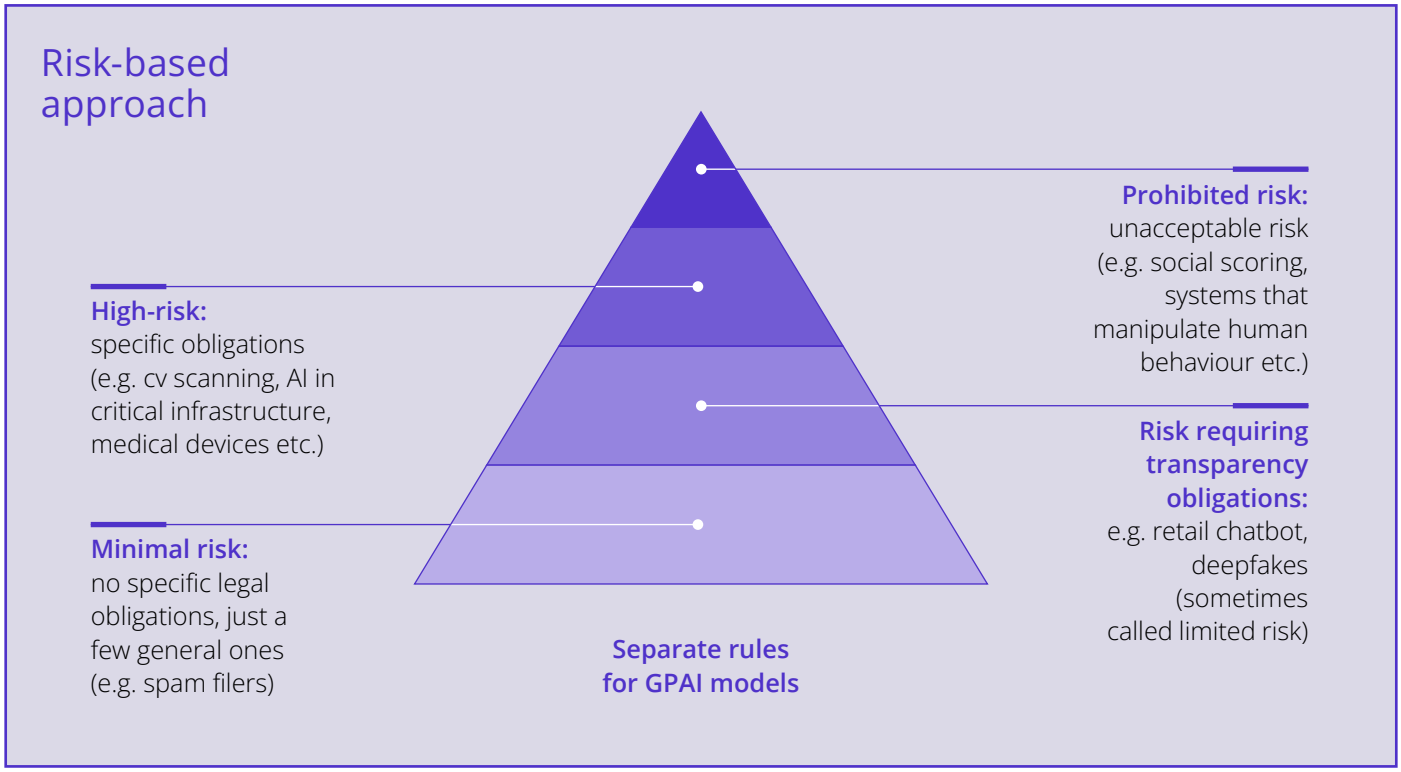

subliminal or deceptive techniques to distort behaviour and cause a harmful decision to be made; creates or expands facial recognition databases; or infers emotions in the workplace or in schools/universities. There are also certain uses relating to biometrics that are prohibited, including real-time biometric ID systems in public (although use is allowed in limited circumstances), and the AI omnibus added AI nudification apps to the list.

(ii) High-risk AI:

These AI Systems are highly regulated. They trigger requirements around risk mitigation, documentation, human oversight, fundamental rights impact assessments and conformity testing. Applicable AI Systems include those:

- intended to be used as safety components in products (or which are themselves products) falling under the EU's product safety legislation (listed in Annex I) and which are required to undergo a third-party conformity assessment before being placed on the market (e.g. toys, aviation, cars, medical devices, lifts etc.), although the AI Omnibus has slightly narrowed the scope (see our blog for more details); or
- which fall within a designated list set out in Annex III of the Act. The list includes: (i) certain (permitted) biometric uses; (ii) critical infrastructure; (iii) education and vocational training; (iv) employment (e.g. in recruitment), worker management and access to self-employment; (v) essential private and public services and benefits (e.g. AI Systems used in accessing healthcare, credit scoring and pricing of life and health insurance); (vi) certain law enforcement uses (e.g. to evaluate the reliability of evidence in investigating or prosecuting criminal offences); (vii) migration, asylum and border control management; and (viii) administration of justice and democratic processes. To trigger the requirements of this category, an AI System must also pose a significant risk of harm to the health, safety or fundamental rights of people. This includes where profiling occurs but does not include systems which carry out narrow procedural tasks, improve the result of a previously completed human activity, detect decision-making patterns or perform certain preparatory tasks.

The Commission published draft guidelines on the classification of high-risk AI systems in May 2026 which contain examples of use cases which are, and are not, high-risk. It also has the right under the Act to add or modify the use cases listed in Annex III (through delegated acts) and to add to or amend the provisions specifying what does not constitute significant risk of harm. (See question 6 for more information).

As mentioned above, high-risk AI Systems which have been on the market or in service before the high-risk rules apply (the dates for which were pushed back in the AI Omnibus to 2 December 2027 for the Annex III high risk use cases and 2 August 2028 for the Annex I) are exempt from the majority of the Act (including the rules relating to high-risk AI Systems) unless significant changes are made to their design. The exemption does not apply to public sector use cases (although there is a longer transition period in such cases) or where an AI System is used on certain large-scale union IT systems. There is no similar exemption for the rules around prohibited AI, which will still apply.

(iii) Transparency Risk AI:

The Act sets out types of AI Systems which, by their nature, present particular transparency risk and therefore require additional transparency and information disclosure obligations (specific to that particular type of AI). These AI Systems are sometimes called "limited risk" AI, although the Act does not use that phrase. Some of the obligations relate to Providers and some relate to Deployers.

The Provider transparency obligations cover:

- **AI Systems intended to interact directly with people (such as chatbots):** Providers must design and develop such systems so that users are informed (unless it is obvious to a reasonably informed person) that they are interacting with an AI System;
- **AI Systems generating synthetic audio, image, video or text content:** Providers must ensure outputs are watermarked or otherwise marked in a machine-readable format and are detectable as artificially generated/manipulated. Providers must also ensure such technology is robust, interoperable and reliable.

The Deployer transparency obligations cover:

- **Emotion recognition or biometric categorisation AI Systems:** Deployers must inform the people exposed about the operation of the AI System and process any data in line with specified EU Regulations and Directives;
- **Deepfakes:** Deployers must disclose that content has been artificially generated or manipulated; and
- **AI Systems generating or manipulating text published to inform the public on matters of public interest:** Deployers must disclose such text as AI generated or manipulated, unless there has been a human review or editorial control and a person (natural/legal) holds editorial responsibility.

The Commission has published both a Transparency Code of Practice, and Transparency guidelines to help organisations comply with the transparency rules.

(iv) Minimal risk AI:

No specific new rules apply if you are using minimal risk AI Systems which perform simple automation tasks without direct interaction with a human (e.g. spam filters). The Commission has stated that "the vast majority" of AI Systems will fall into this category. The Act does, however, impose some general rules which apply to everyone in scope, for example around AI literacy (although these have been amended slightly by the AI Omnibus) and regulatory co-operation.

6. What rules apply to high-risk AI?

High-risk AI Systems are subject to a range of detailed compliance requirements which apply depending on what role, or roles, an organisation plays within the AI value chain.

The majority of the obligations fall on Providers and Deployers, with Providers subject to the most far-reaching obligations. Importers and Distributors are subject to specific checks before placing high-risk AI Systems on the EU market.

Transparency obligations and AI literacy requirements (as discussed above) could also apply, depending on the nature of the AI System involved.

Key technical compliance requirements for high-risk AI include:

- a comprehensive risk management system;
- an obligation to ensure data used to train, validate or test the AI System are relevant, representative and to the best extent possible, free of errors and complete;
- maintenance of thorough technical documentation;
- ensuring the AI System is capable of record-keeping and keeping an event log; and
- obligations as to transparency, human oversight and the accuracy, robustness and cybersecurity of the AI System.

Providers must ensure that their high-risk AI Systems comply with these requirements, and Authorised Representatives, Distributors and Importers also have obligations in relation to them (e.g. to maintain compliance).

Further Provider compliance obligations include:

- AI System registration and record keeping requirements;
- completion of conformity assessments;
- implementation of a quality management system;
- establishment of monitoring systems to gather information from Deployers; and
- cooperation with relevant authorities.

Organisations should be aware that Distributors, Importers, Deployers or other third parties will be considered Providers where they: (i) put their name/trademark on the system (parties can still contractually allocate the attendant obligations); (ii) make a substantial modification to a high-risk AI System and it remains high-risk; or (iii) modify the purpose of a non-high-risk system which makes it high-risk.

Product manufacturers may also be considered Providers if a high-risk AI System is a safety component of a product listed in Section A of Annex I of the Act (including: machinery; toys; lifts etc.) and the high-risk system is (i) placed on the market together with the product under the name or trademark of the product manufacturer; or (ii) put into service under the name or trademark of the Product Manufacturer after the product has been placed on the market.

Compliance obligations for Deployers of high-risk AI Systems include:

- implementation of appropriate technical and organisational measures to ensure use is consistent with instructions;
- appointment of suitably competent individuals in human oversight roles;
- verification of input data to ensure relevance to intended purpose;
- carrying out a fundamental rights impact assessment (for certain public bodies and private entities providing public services e.g. operators of AI Systems which evaluate creditworthiness, or are used in risk assessment for pricing life and health insurance); and
- AI System monitoring, record-keeping and reporting as necessary to Providers, Distributors and authorities.

7. What rules apply to general-purpose AI?

GPAI models are AI models that display significant generality, are capable of competently performing a wide range of distinct tasks and that can be integrated into a variety of downstream systems or applications, as noted in question 2 above. Typically, GPAI models will be trained on large amounts of data using self-supervision at scale. AI models used in R&D and prototyping for pre-market release are excluded. An AI system which is based on a GPAI model and which has the capability to serve a variety of purposes is a GPAI system.

Regulating GPAI models has proved to be one of the most controversial aspects of the Act. When the Act was being agreed, Member states such as Germany and Italy initially opposed stricter regulation for such models and France attempted to block the regulation in its entirety because of the GPAI proposals. Once the Act was in force and the Commission was attempting to agree the GPAI model code of practice (the voluntary guide designed to help organisation comply with the GPAI provisions) there was further disagreement – this time amongst different stakeholder groups.

In terms of the rules themselves, there are rules for both GPAI models generally and GPAI models with systemic risk, both of which will be enforced and monitored by the AI Office at the EU, rather than member state, level.

Providers of GPAI models must:

- draw up and keep updated technical documentation of the model (including on training and testing process and results of evaluation etc). Annex XI sets out a minimum list of technical information to be maintained and to be provided, upon request, to the AI Office and national competent authorities;
- draw up, keep updated and make available information and documentation to Providers of AI Systems that will integrate the GPAI model into their model, enabling such Providers to understand its capabilities and limitations and to comply with their obligations. Annex XII contains a minimum list of such information;
- put a policy in place to comply with EU copyright (and related rights) law, particularly the Directive on Copyright in the Digital Single Market;
- make a summary of the content used for training the GPAI model publicly available (following the AI Office template); and
- cooperate with the Commission and national competent authorities.
- There are additional rules for providers of GPAI models with systemic risk. These include GPAI models that:
- have high impact capabilities (i.e. where cumulative computing power used for training is greater than the agreed threshold – currently FLOPs greater than 10^25); or
- the Commission, or a scientific panel, has decided have equivalent high impact capabilities.

Providers of GPAI models that meet the systemic risk criteria should notify the Commission without delay, and within two weeks of the criteria being met.

Providers may object to their GPAI model being classified as having systemic risks. During the notification process, Providers may present the Commission with arguments to demonstrate that the at-issue GPAI model does not present risk due to its specific characteristics. The Commission will assess and ultimately publish a list of all GPAI models it deems to have systemic risk (without prejudice to the need to respect IP or confidential business information). At this stage, Providers may object, and the Commission may decide to reassess.

For GPAI models that meet the systemic risk threshold, Providers must carry out additional actions to:

- perform model evaluation in line with standardised protocols, including conducting adversarial testing to identify and mitigate risks;
- assess and mitigate possible systemic risks;
- track, document and report (to the AI Office and national competent authorities where relevant) serious incidents and possible corrective measures to address them; and
- ensure adequate levels of cybersecurity (including for the model's physical infrastructure).

8. What are the Act's other measures?

Central to the Commission's aim of fostering innovation, the Act proposes that member states and the European Data Protection Supervisor (on behalf of Union institutions) set up coordinated AI regulatory sandboxes across the EU. These sandboxes promote innovation by providing a controlled environment where Providers can develop, train, test and validate AI Systems for a period before being placed on the market. National competent authorities will be obliged to submit annual reports to the AI Office providing information on the sandbox progress and results which will be available to the public.

The Act also acknowledges the challenges that Small and Medium-sized Enterprises ("SMEs") might face in complying with the new regulations. Some key considerations made for SMEs under the Act include member states offering initiatives such as: priority access to sandboxes, training on how to apply the Act to SMEs, communication channels for advice and responses to queries about implementation, and proportional reductions in fees for Conformity Assessments for certain high-risk AI Systems. Some measures previously reserved for SMEs were also extended to small mid-cap companies under the AI Omnibus.

9. Enforcement and Fines: what happens if you don't comply?

Like many pieces of EU legislation, breach of the Act can lead to large fines. In-scope organisations face fines of (the higher of) €35 million or 7% of global annual turnover in the previous financial year for violations of the prohibited AI practice rules, €15 million or 3% for violations of the Act's obligations (including high-risk compliance obligations, fundamental rights impact assessment and transparency obligations) and €7.5 million or 1.5% for the supply of incorrect information.

Fines for SMEs and start-ups are capped at the lower of the percentages or amounts applicable to each category of violation.

Enforcement of the Act is carried out through a two-layer governance framework at both European Commission level (through the AI Office – supported by a scientific panel of experts – and the AI Board) and at national (member state) level.

For most organisations, enforcement will be carried out at the national level. Each member state will appoint one notifying authority and at least one market surveillance authority – although many Member States have appointed multiple regulators to manage different sectors or risks. National market surveillance authorities will undertake most compliance investigations and enforcement actions.

The AI Office will then enforce the GPAI rules. Changes agreed as part of the AI Omnibus clarify that the AI Office will also oversee certain AI systems built on GPAI models (e.g. where the model and system are developed by the same provider) and those embedded into very large online platforms and very large search engines (as defined under the Digital Service Act). This is, however, subject to certain exceptions where national authorities will remain competent – including in areas like law enforcement, border management, judicial authorities and financial institutions.

This approach aims to ensure harmonised implementation while allowing Member States flexibility to designate competent bodies to carry out effective implementation.

10. Timings: when will the Act apply?

The Act came into force on 1 August 2024, 20 days after being published in the Official Journal of the EU. There was a 2-year transition period, subject to certain exceptions, and some application dates were pushed back following the AI Omnibus.

AI Act Timeline

Key application dates include:

1 August 2024	Act in force
2 February 2025	Prohibited AI practices and general provisions (e.g. regarding AI literacy) apply.
2 August 2025	The rules on GPAI and penalties took effect, and Member States were required to have appointed their notifying authorities and bodies.
2 August 2026	The Act becomes generally applicable.
2 December 2026	The transparency marking rules for providers apply following an additional transition period.
2 August 2027	The GPAI rules will take effect for GPAI placed on the market before 2 August 2025.
2 December 2027	The Annex III high-risk AI System rules will apply. These were originally due to apply in August 2027, but were pushed back following the AI Omnibus.
2 August 2028	The Annex I high-risk AI System rules will apply. These were originally due to apply in August 2027, but were pushed back following the AI Omnibus.
2 August 2030	The high-risk AI System rules will apply to high-risk AI Systems in use by public authorities, where such high-risk AI systems were put on the market before 2 August 2026 (i.e. the Act's general date of application); and
31 December 2030	The Act will apply to AI Systems which are components of the large-scale IT systems established by the legal acts set out in Annex X (which includes a list of EU acts in the areas of freedom, security and justice such as the Schengen Information System).

Practical Steps for Organisations

While many organisations will already be complying with aspects of the Act, the pace of technological and regulatory change means it is important to keep compliance obligations under review.

Organisations should ask themselves:

Do I know what AI I am using?

The first step is to understand what AI you are using in your organisation and whether it falls within the definition of an AI System or GPAI model under the Act.

Am I in scope of the AI Act?

For each AI System, determine:

- Does an exemption apply?
- Are you within the Act's territorial reach?

Even if your organisation does not have an EU base, you may still be in scope if you sell into the EU, or the output of your AI System is used in the EU.

If I am in scope, what role do I play?

For each AI System, determine:

- Are you a Provider, Deployer, Distributor, Importer, Product Manufacturer or Authorised Representative? This will determine the obligations you face.

Which risk category does my AI fall into?

The Act takes a risk-based approach and obligations differ depending on the type of AI you use (e.g. is it prohibited, high, limited or minimal risk or is it a GPAI model?) as well as the role you play (e.g. Provider transparency obligations are different from Deployer transparency obligations).

Am I making changes to an existing AI System?

Modifying an AI System that is already on the market could have substantial consequences under the Act.

For example, in relation to high-risk AI Systems:

- There are certain exemptions which apply for high-risk AI Systems that have been placed on the market before 2 August 2026 (i.e. the general date of application of the Act), but these exemptions fall away where there are significant changes to a high-risk AI System's design.
- If a Distributor, Importer, Deployer or other third party makes a substantial modification to an existing high-risk AI System they may then be treated as a Provider of that system, and subject to the obligations of a Provider under the Act. This will also be the case if they brand the AI System as their own (i.e. put their name or trademark on it).

Changes in use could also bring an AI System into the high risk or prohibited categories, so such use changes should be monitored.

Am I monitoring developments in this space and using the resources available?

There is already a whole host of AI guidance, FAQs, standards and codes of conduct to aid compliance, with more promised.

The Commission has also set up an AI Act Single Information Platform to help organisations understand their obligations (as envisaged in the Act), and shared knowledge gained from those who signed up to the AI Pact (which encouraged early adoption of the Act's measures).

In addition, technology is rapidly evolving in this area, such that organisations will need to monitor advancements as the Act's requirements enter into force to ensure implementation of best practices. Today's best practices might be outdated and insufficient in two years.

Do I need to include AI-specific provisions in my contracts?

While general contractual provisions around compliance with law, performance, liability, IP etc. may provide some protections against AI risk, it is important to keep contracts under review to consider whether you need to include any AI-specific protections in your contracts. For non-EU organisations, protections may include geographic limitations on usage and outputs, among other items. For providers of broadly applicable AI systems who do not want to be caught by the high-risk AI rules, Commission guidelines suggest that any exclusions or limitations on the use (for example, preventing use for high-risk purposes) must be "clearly, concretely and coherently" described across all materials (not just the contractual terms).

Is my AI Governance suitable?

More generally, given AI adoption is advancing at pace, and new AI risks and opportunities are constantly emerging, you should regularly check whether you have appropriate AI governance in place. While AI Act compliance is important – something underpinned by the large, GDPR busting fines provided for under the Act, there are a whole range of AI related risks that are not covered by the Act. A good, regularly reviewed, governance process which helps you set your risk appetite, keep track of your AI use and bring together all relevant stakeholders to identify, manage and monitor the associated risks, can help ensure that AI is developed and/or deployed in your organisation in a responsible way.

Digital Regulation at Slaughter and May

Our Digital Regulation practice covers the full spectrum of digital rules, from AI regulation, competition-focused platform regulation, to online harms and content regulation, and data access and portability. Clients look to us to offer practical, risk-adjusted advice that helps them navigate this ever more complex landscape – and to innovate at speed.

For more information, see our Digital Regulation Practice page or our blog.

Laura Houston
Partner
+44 (0)20 7090 4230
laura.houston@slaughterandmay.com

Will Manley
Head of Digital Regulation
+32 (0)2 737 9431
will.manley@slaughterandmay.com

Natalie Donovan
Head of Knowledge, Tech and Digital
+44 (0)20 7090 4058
natalie.donovan@slaughterandmay.com