

The General-Purpose Artificial Intelligence Code of Practice: what does it mean for me?

On 10 July 2025, the European Commission published the final version of a [Code of Practice](#) on General-Purpose Artificial Intelligence (GPAI).

While the Code is framed as a voluntary tool for GPAI model providers, there are indications in how it has been structured that suggest it could be used as inspiration for a binding regulatory framework. [The Commission has stated](#) that adherence to the Code will help give businesses "a reduced administrative burden" and "increased legal certainty" in terms of compliance with the GPAI-related provisions of the EU AI Act (while not providing a presumption of conformity with the Act, or other relevant laws relating to copyright and data protection). The Code has been supplemented with [additional guidelines](#) on the scope of obligations for GPAI model providers (the Guidance) and is now subject to review by EU Member States and the Commission before it is made final.

In brief

- The development of the Code has not been straightforward with close to 1,000 stakeholders involved and recent calls to "stop the clock" for AI regulation". The Code is now here, but not legally binding – although it could act as the groundwork for a potential, quasi-regulatory framework in the future.
- At the time of writing, certain frontier GPAI model developers, such as [OpenAI](#), [Anthropic](#) and Mistral, have already said they will sign the Code – although there are still some notable omissions. Meta has stated that it will not be signing the Code stating "Europe is heading down the wrong path on AI".
- The Code governs three core legal and regulatory tenets of GPAI development and implementation: [i\) transparency \(including a Model Documentation Form\)](#); [ii\) copyright](#); and [iii\) safety and security](#).
- GPAI model developers will be subject to extensive documentation requirements – and should maintain such documentation for at least 10 years after a model has been withdrawn from the market.
- Businesses implementing GPAI solutions should take note of the Code's requirements and carefully consider their approach to negotiating contracts with GPAI developers and providers accordingly.
- The Guidance supplements the Code with additional guidelines on the scope of obligations for providers of GPAI models. We will report on the Guidance in a follow up article.

Background

The Code was developed through an extensive, iterative process. In covering a broad range of issues ranging from transparency and copyright through to risk assessment and mitigation, it provides a means for GPAI model developers to comply with the AI Act's requirements.

Despite challenges in its development, the Code has now arrived, shortly ahead of the next implementation date of the AI Act's provisions on 2 August 2025. Supplementary guidance on the scope of the EU AI Act was published on 18 July (and we will discuss this in a separate article). You can view a more [detailed timeline of the implementation of the Code and AI Act](#) here.

Are we in scope?

The Code was predominantly formulated to govern providers of commercial GPAI, such as large language model (LLM) developers and text-to-image generation services.

Certain provisions, such as transparency requirements, will not apply to the limited number of open-source and freely licensable GPAI models unless they pose a 'systemic risk'.

However, other stakeholders in the AI value chain should take note of how the Code may affect them. The use of GPAI 'downstream' – in other words, implementation of GPAI within a public-facing offering, such as integrating a GPAI chatbot via an API – raises contractual and liability considerations. For example, Code signatories are required to assist downstream providers in meeting their own obligations under the AI Act by providing technical documentation and related information no later than 14 days upon request. The Code is therefore likely to shape both what downstream developers can expect from their model providers, and the contractual terms they may need to negotiate to ensure regulatory and risk alignment.

Those modifying or fine-tuning existing GPAI models for example, through additional training on proprietary data, or adapting a more generalised model for specific tasks, should also assess how the Code's requirements may apply. While the transparency obligations in the Code principally apply to the original GPAI model developer, Recital 109 of the AI Act makes clear that documentation and disclosure requirements for those fine-tuning a model should be limited to the scope of their modification. In other words, those fine-tuning and adapting existent models should not be expected to disclose or take responsibility for aspects of the base model above and beyond that which they did not change or influence.

Finally, it is also important to note the extra-territorial effect of the AI Act, which applies not only to organisations established in the EU, but also to those placing AI systems on the EU market or using them within the EU, regardless of where the provider is based. In practice, this means that even non-EU companies integrating or adapting GPAI models may find themselves subject to EU regulatory expectations, and indirectly, to the standards set by this Code. The Commission is therefore, in practice, positioning the Code and the Guidance as global benchmarks, and suppliers operating internationally should not assume that they are automatically out of scope.

Is the Code binding?

The Code is, at this stage, voluntary and non-binding, and is intended to assist GPAI providers to comply with their obligations under the AI Act. It does not supersede or otherwise affect existing requirements to comply with other applicable laws, directives, and regulatory frameworks, such as copyright and data protection laws. While it provides a means of demonstrating compliance with the AI Act, it does not provide a presumption of conformity, and, indeed, conformity can be achieved through other means. The Commission has also [confirmed](#) that providers will be able to benefit in effect from a grace period to demonstrate "good faith" adherence to the Code's requirements following their signature of the Code.

However, wording in the Code's final version does signal a marked change in approach. Notably, the final draft has removed language that established compliance thresholds tied to, for example, 'reasonable endeavours' and instead adopts more directive and prescriptive obligations. For now, compliance with the Code is heavily encouraged as a means to demonstrate compliance with the AI Act's obligations, [but in the Commission's own words](#), "*from 2 August 2025 onwards, the Commission will enforce full compliance with all obligations for providers of general-purpose AI models with fines*".

What does the Code say?

The Code is split into three distinct chapters:

Transparency

The Code's rather descriptive [transparency provisions](#) require GPAI developers to maintain technical information regarding several factors concerning a model's development, proportional to the model's size. The bulk of these requirements are listed in the Commission's 'Model Documentation Form', which covers the following:

- Basic information, such as the name and size of the model, the name of its provider, and the date of release
- A general description of the model's architecture, including whether the model is a fine-tuned derivative of a prior GPAI model
- Which input and output formats (modalities) are supported
- The licence terms on which the model has been made available
- Details of applicable 'acceptable use' policies
- A general description of how the model was trained, including a breakdown of each training stage, key technical/design choices and assumptions, as well as how its training data was sourced
- A measurement (or estimate) of the amount of energy used in the model's development

While Code signatories are required to continuously maintain this information for the lifetime of a model, and for an additional ten years after it has been removed from the market, there is no requirement that it be made publicly available (though signatories are encouraged to consider whether it can be in order to promote public transparency). The wording of this chapter suggests that, instead, such information should be collated primarily for if and when it is later requested by a third party, such as a regulatory authority, or downstream providers, subject to safeguards with respect to confidentiality and IP.

Copyright

This is the Code's shortest and, at least on its face, most 'straightforward' [chapter \(despite the interplay between copyright and GenAI being the subject of a significant number of ongoing global disputes and government consultations\)](#). The various commitments in this chapter reflect the obligations on model providers in Article 53 of the AI Act.

- Code signatories must draw up, keep up-to-date, and implement a copyright policy, which they are encouraged to make publicly available.
- Providers which crawl the web and scrape data should only reproduce and extract lawfully accessible copyright-protected content. They should not circumvent measures such as paywalls and not scrape websites that are recognised repeat infringers (a list of such sites will be published).
- They must also identify and respect machine-readable reservations of rights, such as Robot Exclusion Protocol (robots.txt) files. While the Code also suggests that developers should identify and comply with other appropriate machine-readable protocols that may be adopted, it leaves development and agreement of those mechanisms to good faith discussions between rightsholders and AI providers. Currently, this risks being a problematic approach. It is well-known that discussions between these two groups have been difficult. It is worth noting here

that the UK Government has recently [convened expert working groups](#) comprising representatives of the creative industries and AI sectors to try and develop 'practical, workable solutions' to outstanding rights reservation issues. The market has also begun to explore novel approaches to disincentivising scraping practices, such as Cloudflare blocking scrapers from accessing a site unless they pay to do so. Businesses reliant upon readership and page visits are yearning for these solutions to bear fruit, as evidence suggests that simply 'turning off' scraping access entirely has a detrimental impact on a website's SEO.

- Of note, Code signatories that also provide an online search engine must ensure compliance with a rights reservation does not "directly lead to adverse effects on the indexing of the content, domain(s) and/or URL(s), for which a rights reservation has been expressed, in their search engine".
- Code signatories must develop 'appropriate and proportionate' technical measures to mitigate the risk of models generating outputs that reproduce protected, copyrighted works.
- They must also implement and maintain a process, with a listed 'point of contact', for rightsholders to complain about non-compliance with the Code, and to deal with those complaints.

Safety and security

The Code's [safety and security chapter](#) is by far the most detailed and is primarily intended to help businesses comply with Article 55 of the AI Act. At a high level, Code signatories must develop and continuously revise a sufficient governance framework for safety and security, including a process to identify, analyse, and mitigate or accept 'systemic risk' as defined in the AI Act.

The Code also tasks GPAI developers with advancing the state of the art of AI safety and security while not unduly reducing the beneficial capabilities of the models they develop. It emphasises achieving "equal or superior safety or security outcomes", which are "recognised as advancing the state of the art in AI safety and security and meriting consideration for wider adoption". Less onerous and more proportionate reporting requirements will be applied to SMEs and SMCs, in line with Principle (F) of the Code's Recitals. However, these businesses can elect to voluntarily adhere to the more stringent requirements if they wish. This overarching principle aims to avoid placing disproportionate burdens on smaller actors, particularly those without the scale or resources of major GPAI developers.

Notably, the AI Office has already provided some further clarifications in the Guidance on the relevant thresholds at which a GPAI model will be deemed to have systemic risk. This is one of several examples where clarificatory guidance and statutory amendments will be likely to have significant implications for AI regulation and the need for such guidance to keep pace with the rapid development of AI. For now, Article 51(2) of the AI Act states that an indicator that a model may be of systemic risk is when its level of compute power (when measured in floating point operations) "is greater than 10^{25} " – a threshold that almost all frontier models now readily surpass.

Key takeaways

It remains to be seen whether the Code will lead to a marked change in AI market practices and governance in the coming months. Even with the Commission's subsequent release of the Guidance, there is a possibility that developers will implement divergent compliance measures – if they agree to be a signatory to the Code at all. Despite such challenges, the Commission's Executive Vice-President for Tech Sovereignty, Security and Democracy Henna Virkkunen, has called for organisations to begin compliance efforts sooner rather than later:

"I invite all general-purpose AI model providers to adhere to the Code. Doing so will secure them a clear, collaborative route to compliance with the EU's AI Act."

Authors: Harry Clark, Associate ([Harry Clark | Associate | Innovation](#)) and Anne Rose, Managing Associate ([Anne Rose | Managing Associate | Innovation](#)) at Mishcon de Reya LLP