

AI Act duties for general-purpose AI model providers start to apply

Source copy kept by wheresthe.ai for the story "AI Act duties for general-purpose AI model providers start to apply" (<https://wheresthe.ai/d/eu-gpai-obligations-apply-2025-08/>).

Source: <https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: a78dec941f20d6a53a04dbf633e5c1d3db6d19b0f1db215e3fce12b9d32861fd

Questions and answers on commission's guidelines on the scope of the obligations for providers of General-Purpose AI (GPAI) models.

General FAQ

Why are general-purpose AI models specifically addressed in the AI Act? General-purpose AI models play a significant role in AI innovation because they can be used for various tasks and integrated into many downstream AI systems. This places particular responsibilities on their providers. In particular, they must make available information to providers of AI systems who intend to integrate the model into their AI systems, to allow these downstream providers to understand the capabilities and limitations of the model, and to fulfil their own obligations under the AI Act. Since these models are usually trained on vast amounts of data that may include copyright-protected content, the AI Act requires providers to establish copyright policies, and publish summaries of the content used to train their general-purpose AI models. Finally, general-purpose AI models may present systemic risks that can have a significant impact on the Union market. Providers of such models are subject to additional obligations aimed at assessing and mitigating these systemic risks. These obligations include carrying out model evaluations, incident reporting, and ensuring adequate cybersecurity protections.

What is the purpose of these guidelines, and how do they relate to the AI Act? These guidelines clarify how the Commission interprets key concepts in the AI Act, specifically in relation to the obligations for providers of general-purpose AI models that enter into application on 2 August 2025. In doing so, the guidelines provide legal certainty to providers. They help actors along the AI value chain determine whether their model qualifies as a general-purpose AI model, whether they are the provider placing it on the market, whether they qualify for exemptions, and what to expect regarding the Commission's enforcement of the obligations in the AI Act. These guidelines apply solely to the AI Act and not to other Union laws.

What legal obligations do providers of general-purpose AI models face under the AI Act as of 2 August 2025? Providers of general-purpose AI models must: draw up and maintain technical documentation about the model, including details of the development process, to provide to the AI Office upon request. National competent authorities can also ask the AI Office to request information on their behalf when this information is needed for their supervisory tasks; provide information and documentation to downstream AI system providers to help them understand the model's capabilities and limitations and comply with their own obligations; implement a policy to comply with Union copyright law and related rights, using state-of-the-art technologies to identify and respect rights reservations; publish a sufficiently detailed summary of the content used for training the model; if established outside the EU, appoint an authorised representative in the Union before placing their model on the market. Providers of general-purpose AI models released under a free and open-source license may be exempt from the first obligations under certain conditions. However, providers of general-purpose AI models with systemic risk, including open-source models, face additional obligations. These providers must, for instance, notify the Commission when developing a model with systemic risk and take steps to ensure the model's safety and security. Providers can demonstrate compliance through a Code of Practice assessed as adequate, or via alternative adequate means.

How do the guidelines help identify whether a model qualifies as a general-purpose AI model? The guidelines help identify whether a model qualifies as a general-purpose AI model if the computational resources used for its training (training compute) exceed 10^{23} floating point operations (FLOP) and it can generate language (text or audio), text-to-image, or text-to-video. This compute threshold corresponds to what is typically used to train models with one billion parameters on large datasets. Moreover, the chosen modalities enable flexible content generation capable of a wide range of distinct tasks. However, this is not an absolute rule - models meeting this criterion may exceptionally not qualify as general-purpose AI models if they lack significant generality, while models below this threshold may still be general-purpose AI models if they display significant generality and can

competently perform a wide range of tasks.

How can AI companies know whether they must comply with obligations for general-purpose AI models under the AI Act? AI companies must comply with the obligations for providers of general-purpose AI models if they are providers placing such a model on the Union market. Qualifying as a general-purpose AI model : Based on the indicative criterion given in the guidelines, models trained with an amount of compute that exceeds 10^{23} FLOP and that can generate language, text-to-image, or text-to-video are considered general-purpose AI models, though there may be exceptions. Being the “provider”: A “provider” is the entity that develops or has a general-purpose AI model developed and places it on the market under its own name or trademark. This includes companies established outside the EU that place models on the Union market. “Placing on the market”: This means the first availability on the Union market, whether via APIs, downloads, cloud services, integration into applications, or other means. This applies even when models are integrated into AI systems, or used for internal processes essential for providing a product or service to third parties, or affecting the rights of natural persons in the Union. The guidelines provide examples to identify the provider in various scenarios, including collaborative development and third-party arrangements, and clarify what actions constitute placing on the market.

How does the AI Act balance innovation and regulation for developers modifying existing general-purpose AI models? The AI Act balances innovation and regulation by recognising that not every modification or fine-tuning of a general-purpose AI model should be treated as creating a new model. As such, actors modifying or fine-tuning a model are not automatically subject to all the obligations for providers of general-purpose AI models. The guidelines clarify that these actors become providers only in exceptional circumstances, specifically when the modification or fine-tuning uses more than one-third of the original model's training compute. This high threshold means most fine-tuning, adaptations, and minor modifications will not subject developers to the obligations for providers. The Commission, in support of innovation, further limits obligations for significant modifications to documenting the modification itself — the changes made and new training data used — rather than requiring full documentation of the entire model. This approach ensures that the vast majority of developers can innovate by building on existing models without excessive regulatory burden.

Are these guidelines legally binding? These guidelines are not legally binding. An authoritative interpretation of the AI Act may only be given by the Court of Justice of the European Union. Nevertheless, these guidelines set out the Commission's interpretation and application of the AI Act, on which it will base its enforcement action. This helps providers comply with their obligations and supports the effective implementation of the AI Act. Nonetheless, a case-by-case assessment will always be necessary to consider the specifics of each individual case.

When is a general-purpose AI model a general-purpose AI model with systemic risk? A general-purpose AI model is classified as having systemic risk if it meets one of two conditions. Compute threshold condition: The model has capabilities that match or exceed those of the most advanced models. The AI Act presumes this for models trained with a cumulative amount of computational resources exceeding 10^{25} FLOP (the ‘compute threshold’) to have such capabilities. This threshold suggests that the model could significantly impact the Union market due to its reach or potential negative effects on public health, safety, security, fundamental rights, or society. The Commission must adjust this threshold when necessary to account for technological developments. Designation condition : The Commission can designate a model as a general-purpose AI model with systemic risk either on its own initiative or in response to a qualified alert from the scientific panel if the model's capabilities or impact are equivalent to those of the most advanced models. This provision accounts for models that may pose systemic risks even if they do not meet the compute threshold. Once a model meets either of these conditions, its provider must comply with additional obligations, including assessing and mitigating systemic risks.

If a general-purpose AI model meets the compute threshold, does it automatically become a general-purpose AI model with systemic risk? When a general-purpose AI model meets the compute threshold, or it becomes known that this threshold will be met, the provider must notify the Commission without delay, and within two weeks at the latest. In this notification, the provider may present arguments to demonstrate that, despite meeting the threshold, the model does not have capabilities matching or exceeding those of the most advanced models, or that it does not present systemic risks for other reasons. Therefore, it should not be classified as a general-purpose AI model with systemic risk. Providers may specifically argue that their model does not present systemic risks because it lacks high-impact capabilities - those that match or exceed the capabilities of the most advanced models, and systemic risks associated with such high-impact capabilities. The Commission will assess the arguments submitted by the provider and decide whether to accept or reject them. If the Commission accepts the arguments, the model will no longer be classified as a general-purpose AI model with systemic risk, and its provider will not be subject to the related obligations from the moment they are informed of the acceptance

decision. If the Commission rejects the arguments, the model will be confirmed as a general-purpose AI model with systemic risk, and the provider will be subject to the obligations for such models from the moment the model meets the compute threshold.

Are all general-purpose AI models subject to the same requirements, or are there exemptions — such as for open-source models? Not all providers of general-purpose AI models are subject to the same requirements, as there are exemptions for providers of models released under free and open-source licenses. Specifically, providers of such models may be exempt from the requirement to: maintain technical documentation for authorities; provide documentation to downstream AI system providers; appoint an EU representative (for non-EU providers). These exemptions only apply if the model: is released under a truly free and open-source license that allows access, use, modification, and distribution without monetisation; has parameters, including weights, architecture, and usage information publicly available; is not classified as a general-purpose AI model with systemic risk, as providers of those models must comply with all obligations regardless of whether the model is open-source. The exemptions acknowledge that open-source models support research and innovation while providing transparency through their open nature. However, providers of open-source models are still required to meet the copyright policy obligation and publish a training data summary.

What are the additional requirements for general-purpose AI models with systemic risk? In addition to standard obligations, providers of all general-purpose AI models, and providers of general-purpose AI models with systemic risk must: Model evaluation: Perform model evaluation using standardised protocols and state-of-the-art tools, including conducting and documenting adversarial testing to identify and mitigate systemic risks. Risk assessment: Assess and mitigate possible systemic risks at Union level, including their sources, which may stem from the development, placing on market, or use of these models. Incident reporting: Track, document, and report relevant information about serious incidents and possible corrective measures to the AI Office and, as appropriate, national authorities without undue delay. Cybersecurity safeguards: Ensure adequate cybersecurity protection for both the model and its physical infrastructure, to prevent theft, misuse or widespread consequences of malfunction. Providers can demonstrate compliance through adherence to the General-Purpose AI Code of Practice or by showing alternative adequate means of compliance. If providers choose the latter, they must present arguments for why such means are adequate, which will be assessed by the Commission.

How do the Commission guidelines on the scope of the obligations for providers of general-purpose AI models complement the General-Purpose AI Code of Practice released last week? The guidelines complement the Code of Practice by clarifying key concepts in the AI Act, including how the Code can be used to demonstrate compliance. While the guidelines provide an interpretative framework for understanding the obligations of providers of general-purpose AI models, the Code offers specific measures that providers can implement to demonstrate that they meet these obligations.

How were the guidelines developed, and to what extent were AI model providers and other stakeholders consulted? The Commission developed these guidelines through an inclusive consultation process to ensure they reflect practical experience and diverse perspectives. A public multi-stakeholder consultation was carried out from 22 April to 22 May 2025, inviting input from general-purpose AI model providers, downstream providers, civil society, academia, experts, public authorities, and other stakeholders. The guidelines also incorporated feedback from Member States through the European AI Board and drew on expertise from the Commission's Joint Research Centre pool of experts.

What kind of documentation obligations apply to providers of general-purpose AI models under the AI Act? Providers of general-purpose AI models must fulfil two main documentation obligations under the AI Act: Technical documentation for authorities: Providers must draw up and maintain comprehensive technical documentation for authorities. This includes information about the model's architecture, training process, training, testing, and validation data, computational resources, and energy consumption. This documentation must be made available to the AI Office upon request, which may also act on behalf of national competent authorities. Documentation for downstream providers: Providers must also create and maintain separate documentation – those who intend to integrate the general-purpose AI model concerned into AI systems ('downstream providers'). This documentation should include general information about the model's intended tasks, technical integration requirements, input/output specifications, and training data. It must be proactively made available to downstream providers to enable them to understand the model's capabilities and limitations and comply with their obligations under the AI Act. Both sets of documentation must be kept up to date throughout the model's lifecycle. General-purpose AI models released under a free and open-source license may be exempt from these documentation requirements, provided they meet certain conditions about additional information that must be made available.

However, this exemption does not apply to general-purpose AI models with systemic risk.

What are the public transparency and copyright-related obligations for all providers of general-purpose AI models? Providers of general-purpose AI models must fulfil two main public transparency and copyright-related obligations under the AI Act: Copyright policy: Providers must put in place a policy to comply with Union copyright law and related rights. This includes identifying and respecting rights reservations under Union copyright law. Public summary of training content: Providers must make a sufficiently detailed summary of the content used for training their models publicly available. These obligations apply to all providers of general-purpose AI models, including providers of open-source models. The General-Purpose AI Code of Practice provides detailed guidance on how providers can meet the copyright compliance obligation. For the public summary requirement, the Commission is developing a template and accompanying guidelines that providers will need to use for presenting the public summary.

How does the AI Act ensure downstream developers have sufficient information to use general-purpose AI models responsibly? The AI Act ensures downstream developers have sufficient information through mandatory documentation requirements for providers of general-purpose AI models. Providers must create and actively provide documentation designed specifically for downstream AI system providers, which includes: the model's intended tasks and acceptable use policies; technical specifications including architecture, parameters, input/output modalities and format; integration requirements such as instructions for use, infrastructure needs, and necessary tools; information about training, testing, and validation data type, provenance, and curation methodologies. By mandating this information sharing, the Act creates a transparency framework that supports responsible innovation across the entire AI value chain.

What are the next steps on enforcement? The AI Act obligations for providers of general-purpose AI models enter into application on 2 August 2025. From this date, providers placing these models on the market must comply with their AI Act obligations and notify the AI Office without delay about models with systemic risk to be placed on the EU market. In the first year from 2 August 2025 onwards, the AI Office will offer to collaborate closely in particular with providers who adhere to the Code of Practice to ensure that models can be placed on the EU market without delay. If providers adhering to the Code do not fully implement all commitments immediately, the AI Office will not consider them to have broken their commitments under the Code. Instead, the AI Office will consider them to act in good faith and will be ready to collaborate to ensure full compliance. However, from 2 August 2026 onwards, the Commission will enforce full compliance with all obligations for providers of general-purpose AI models, including through fines. Providers of general-purpose AI models placed on the market before 2 August 2025 must comply with the AI Act obligations by 2 August 2027.