



General-purpose AI models

Articles 51 to 56, and the two annexes they point at: what every model provider owes, what systemic risk adds, and what an open-source release does not lift.

5

duties on every model

5

more for systemic risk

2

routes to systemic risk

2

documentation annexes

The model, not the system built on it

Chapter V attaches to the model layer, so none of the risk tiers in Article 6 decide anything here. One division does: every general-purpose model carries 5 duties, and a model classified as carrying systemic risk carries 5 more.

Keep technical documentation of the model

Providers must draw up and keep up to date technical documentation of the model, covering its training and testing process and the results of its evaluation, and hand it to the AI Office or a national competent authority on request.

Exemptions. Models released under a free and open-source licence meeting the Article 53(2) conditions — but not if the model carries systemic risk.

Article 53(1)(a)

Give downstream providers what they need to comply

Providers must make information and documentation available to providers of AI systems that intend to integrate the model, sufficient for them to understand its capabilities and limitations and to meet their own obligations under the Regulation.

Exemptions. The same Article 53(2) open-source carve-out, again excluding systemic-risk models. Intellectual property rights, confidential business information and trade secrets are preserved.

Article 53(1)(b)

Put a copyright policy in place

Providers must have a policy to comply with Union copyright law, and in particular to identify and respect reservations of rights expressed under Article 4(3) of the Copyright in the Digital Single Market Directive, using state-of-the-art technologies.

Exemptions. No exemption: the Article 53(2) carve-out reaches points (a) and (b) only.

Article 53(1)(c)

Publish a summary of the training content

Providers must draw up and make publicly available a sufficiently detailed summary of the content used to train the model, following the template provided by the AI Office.

Exemptions. No exemption: like the copyright policy, this sits outside the Article 53(2) carve-out.

Article 53(1)(d)

Appoint an authorised representative in the Union

A provider established outside the Union must appoint an authorised representative established in the Union, by written mandate, before placing a general-purpose AI model on the Union market.

Exemptions. Providers already established in the Union.

Article 54

Two ways in, one of them a number

Article 51(2) supplies a presumption, not a definition. The Commission can designate a model on the Annex XIII criteria with no reference to compute at all — and reach alone is one of those criteria.

High impact capabilities

A model has high impact capabilities when they are evaluated as such on the basis of appropriate technical tools and methodologies, including indicators and benchmarks. Article 51(2) adds a presumption: a model is presumed to have them when the cumulative compute used for its training exceeds 10^{25} floating point operations.

Article 51(1)(a) with 51(2)

A decision of the Commission

The Commission may decide, on its own initiative or following a qualified alert from the scientific panel, that a model has capabilities or an impact equivalent to high impact capabilities, judged against the Annex XIII criteria. No compute threshold is involved in this route.

Article 51(1)(b) with Annex XIII

Notify the Commission within two weeks

Where a model meets the high-impact-capability condition in Article 51(1)(a), the provider must notify the Commission without delay, and in any event within two weeks of the requirement being met or of it becoming known that it will be met.

Exemptions. A provider may argue in the notification that the model exceptionally does not present systemic risk; until the Commission accepts that, the model is classified as carrying systemic risk.

Article 52(1)

Evaluate the model, adversarially and on the record

Providers of systemic-risk models must perform model evaluation using standardised protocols and tools reflecting the state of the art, including conducting and documenting adversarial testing aimed at identifying and mitigating systemic risks.

Article 55(1)(a)

Assess and mitigate systemic risk at Union level

Providers must assess and mitigate possible systemic risks at Union level, including their sources, arising from the development, the placing on the market or the use of the model.

Article 55(1)(b)

Report serious incidents to the AI Office

Providers must keep track of, document and report serious incidents and possible corrective measures to the AI Office, and as appropriate to national competent authorities, without undue delay.

Article 55(1)(c)

Protect the model and the infrastructure under it

Providers must ensure an adequate level of cybersecurity protection for the model and for the physical infrastructure of the model.

Article 55(1)(d)

What Article 53(2) actually lifts

The most misquoted provision in the chapter. It lifts two duties, not four, and only where every one of its conditions is met.

Lifted

53(1)(a) technical documentation, and 53(1)(b) downstream information.

Article 53(2)

Still owed

53(1)(c) the copyright policy, and 53(1)(d) the public training-content summary.

Article 53(1)(c) and (d)

Unavailable to a systemic-risk model

The exception does not apply at all to a model classified as carrying systemic risk. Recital 104 is explicit that openness is not on its own a reason to drop those obligations.

Recital 104

The conditions, all of them

- Released under a free and open-source licence allowing access, usage, modification and distribution
- Parameters, including the weights, made publicly available
- Information on the model architecture made publicly available
- Information on model usage made publicly available

Article 53(2)

What the two annexes ask for

Article 53(1) points (a) and (b) each name an annex, and the two are not the same document. One goes to regulators on request; the other goes to the companies building on your model.

Annex XI — for the ai office and national competent authorities, on request

- A general description: intended tasks, the systems it can be integrated into, acceptable use policies, release date and distribution methods, architecture and parameter count, input and output modality and format, and the licence
- Design specifications and training methodology, including key design choices with their rationale, and what the model is designed to optimise for
- Data used for training, testing and validation: type and provenance, curation methodology, number of data points, scope, how it was obtained and selected, and the measures used to detect unsuitable sources and identifiable biases
- Computational resources used to train the model, training time, and known or estimated energy consumption

Article 53(1)(a)

Annex XII — for downstream providers integrating the model into their systems

- A general description: intended tasks, the systems it can be integrated into, acceptable use policies, release date and distribution methods, interaction with hardware or software outside the model, relevant software versions, architecture and parameter count, modality and format, and the licence
- The technical means required to integrate the model into an AI system
- Input and output modality and format and their maximum size, such as context window length
- Information on the data used for training, testing and validation: type, provenance and curation methodology

Article 53(1)(b)

Five readings that cost money

“Under 10²⁵ FLOP means no systemic-risk duties.”

10²⁵ is a presumption under Article 51(2), not the definition. Article 51(1)(b) lets the Commission designate a model on the Annex XIII criteria with no reference to compute at all, and reach alone — 10 000 registered business users in the Union — is one of those criteria.

[Article 51\(1\)\(b\), Annex XIII](#)

“Our model is open source, so Chapter V does not apply.”

Article 53(2) lifts points (a) and (b) only, and only where the licence and the published parameters, architecture and usage information all meet its conditions. The copyright policy and the public training-content summary still apply, and the whole carve-out falls away for a systemic-risk model.

[Article 53\(2\)](#)

“We fine-tuned someone else's model, so we inherit all their obligations.”

Recital 109 limits the obligations in the case of a modification or fine-tuning to that modification or fine-tuning, for example by complementing the existing technical documentation rather than recreating it.

[Recital 109](#)

“We only call an API, so the model duties are ours too.”

Chapter V binds the provider of the model. Integrating a model into your own system makes you a downstream provider with system-level duties; Article 53(1)(b) exists precisely so that the model provider hands you what you need for them.

[Articles 53\(1\)\(b\), 101](#)

“A code of practice is optional paperwork.”

It is the route the Act names. Until a harmonised standard is published, providers may rely on a code of practice to demonstrate compliance; a systemic-risk provider who adheres to neither must demonstrate alternative adequate means of compliance for assessment by the Commission.

[Articles 53\(4\), 55\(2\), 56](#)

You provide a model, or you build on one

HOLISTIC AI

The enterprise AI governance platform

Holistic AI governs AI across the lifecycle: an inventory of every system, the classification behind each one, the evidence a supervisory authority will ask for, and continuous monitoring once it is live.

[EU AI Act readiness assessment →](https://holisticai.com/eu-ai-act-readiness)

holisticai.com/eu-ai-act-readiness

[Get a demo →](https://holisticai.com/demo)

holisticai.com/demo

Classify one system

The risk calculator walks the same sequence the Act does and names the provision behind each answer.

euaiact.com/eu-ai-act-risk-calculator

Check your exposure

The readiness check scores the whole estate against the obligations that already bind and the ones that are coming.

euaiact.com/eu-ai-act-readiness



The EU AI Act, as a working tool

euaiact.com is maintained by Holistic AI as an independent, source-verified reference: the complete legal text with the Omnibus amendments marked, a scope finder, a risk classifier, eight sector guides and a 31-step compliance checklist — all free.

YOUR NEXT STEP

Find out where you actually stand

Classify one system in a few minutes, or have Holistic AI assess the whole portfolio and produce the evidence behind it.

[Classify a system →](#)

[Start the assessment →](#)

euaiact.com/eu-ai-act-risk-calculator

holisticai.com/eu-ai-act-readiness

Or book a walkthrough of the platform: holisticai.com/demo