The 30 AI Act questions · Written by Yobel Tzegai · Last checked 2 November 2026

The 30 AI Act questions your enterprise customers will ask

The kinds of AI Act questions that show up in enterprise vendor questionnaires and audits, grouped by topic, with what a good answer includes. Compiled by Complipath (complipath.io), and free as a page or a PDF, with no form.

Get my risk classes, free

Book a 20-min call

20 minutes with the founder.

Download the 30 AI Act questions

Download the 30 AI Act questions as a PDF: the same questions, in the same order, with nothing to fill in first. No form and no e-mail address.

Questionnaire waiting? We answer it in 48 working hours.

Inventory and roles

- Which AI systems are in the service you sell us? What a good answer includes: A named list, one line per system, saying what each does and why it meets the definition of an AI system in Article 3, point (1). Features built on a third party's model belong on the list too. (Article 3, point (1)). Read: AI inventory: the register that dates its own law.
- For each system, are you the provider or the deployer? What a good answer includes: A role per system: provider where you develop it or have it developed and place it on the market or put it into service under your own name or trademark, paid or free (Article 3, point (3)); deployer where you use it under your authority, other than in a personal non-professional activity (Article 3, point (4)). One company can hold both roles for different systems. (Article 3, points (3) and (4)). Read: Am I a provider or a deployer?.

- Have you rebranded, substantially modified or repurposed someone else's AI system? What a good answer includes: Yes or no per system, with the reason. Under Article 25(1) you become the provider of a high-risk system if you put your name or trademark on one already on the market, without prejudice to contracts allocating the obligations otherwise; make a substantial modification after which it remains high-risk; or change the intended purpose of a system not classified as high-risk so that it becomes high-risk. (Article 25(1)). Read: When does a deployer become a provider?.
- Which general-purpose AI models do you build on, and what do their providers tell you? What a good answer includes: The models per system, and the information and documentation their providers make available to you under Article 53(1), point (b), so that you understand each model's capabilities and limitations, with intellectual property and trade secrets protected. Models released under a free and open-source licence that meet Article 53(2) are exempt from that duty, unless they are general-purpose AI models with systemic risk. (Article 53(1), point (b); Article 53(2)). Read: What are the general-purpose AI (GPAI) obligations under the EU AI Act?.
- How do you build AI literacy in the people who run and use your AI? What a good answer includes: The measures you take for staff and others operating or using AI on your behalf, matched to their technical knowledge, experience, education and training, the context of use and the people the systems are used on (Article 4(1), as replaced by Regulation (EU) 2026/1744 (the Digital Omnibus on AI)). The duty is to take measures that support AI literacy, not to guarantee any individual's level. (Article 4(1), as replaced by Regulation (EU) 2026/1744). Read: Who does the EU AI Act's AI literacy duty apply to?.

Risk classification

- Is any AI system you provide or use high-risk, and by which route? What a good answer includes: A classification per system with the route it rests on: Article 6(1) for a safety component of, or itself, a product covered by Annex I legislation that must undergo third-party conformity assessment; Article 6(2) for a system listed in Annex III, subject to Article 6(3). Where neither applies, say why. (Article 6(1) and (2)). Read: How to classify your AI system.
- If you decided an Annex III system is not high-risk, where is that assessment? What a good answer includes: The assessment, documented before the system was placed on the market or put into service (Article 6(4)), naming the Article 6(3) condition it rests on: a narrow procedural task, improving a previously completed human activity, detecting decision-making patterns or deviations without replacing or influencing the previous human assessment without proper human review, or a preparatory task. A system that profiles natural persons is always high-risk, and a provider relying on the exemption registers itself and the system

under Article 49(2). (Article 6(3) and (4); Article 49(2)). Read: When does the Article 6(3) exemption apply?.

- Do you use AI for any practice the AI Act prohibits? What a good answer includes: An answer per system, checked against every point of Article 5(1), first subparagraph. Those prohibitions have applied since 2 February 2025, except points (ba) and (bb), which Regulation (EU) 2026/1744 inserted and which apply, with Article 5(1a) and (1b), from 2 December 2026 (Article 113, third paragraph, point (a), as replaced). (Article 5(1), first subparagraph; Article 113, third paragraph, point (a), as replaced by Regulation (EU) 2026/1744). Read: The ten prohibited AI practices in Article 5.
- Is any of your AI a safety component of a regulated product, or a regulated product itself? What a good answer includes: Per system, the Annex I act it falls under (if any) and whether that act requires third-party conformity assessment of the product, because Article 6(1) needs both. Since Regulation (EU) 2026/1744, a system solely used for non-safety related aspects of user assistance, performance optimisation, service efficiency, automation or convenience or quality control is not a safety component, unless its failure or malfunctioning would endanger health and safety (Article 6(1a) and (1b)). (Article 6(1); Article 6(1a) and (1b), as inserted by Regulation (EU) 2026/1744; Annex I). Read: Annex I vs Annex III: two routes to high-risk.
- From when do the high-risk obligations apply to your systems? What a good answer includes: A date per high-risk system: Chapter III, Sections 1, 2 and 3, except Article 6(5), apply from 2 December 2027 for Annex III systems and from 2 August 2028 for Annex I systems (Article 113, third paragraph, point (c), as replaced by Regulation (EU) 2026/1744). Registration under Article 49 sits in Section 5, which keeps the general date of 2 August 2026 (Article 113, second paragraph); the amending act does not say how that works before the classification rules apply. (Article 113, third paragraph, point (c), as replaced by Regulation (EU) 2026/1744). Read: High-risk obligations from 2 December 2027.

Transparency

- Do you tell people when they are interacting with an AI system? What a good answer includes: As provider, yes for each system intended to interact directly with natural persons, unless that is obvious to a reasonably well-informed, observant and circumspect person or the law-enforcement exception applies (Article 50(1)). Say where the notice appears: Article 50(5) requires it in a clear and distinguishable manner at the latest at the first interaction, meeting the applicable accessibility requirements. (Article 50(1) and (5)). Read: Chatbots and voice agents under the EU AI Act.
- Is content your AI generates marked as AI-generated? What a good answer includes: For systems generating synthetic audio, image, video or text, outputs marked in a machine-

readable format and detectable as artificially generated or manipulated (Article 50(2)), unless the system performs an assistive function for standard editing, does not substantially alter the input data or its semantics, or the law-enforcement exception applies. Systems placed on the market before 2 August 2026 have until 2 December 2026 (Article 111(4)). (Article 50(2); Article 111(4), as added by Regulation (EU) 2026/1744). Read: Article 50: the four limited-risk disclosures.

- Do you disclose deep fakes and AI-generated text published to inform the public? What a good answer includes: Where you deploy such a system, yes: a deep fake is disclosed as artificially generated or manipulated, and for an evidently artistic, creative, satirical, fictional or analogous work the duty is limited to disclosing that such content exists, without hampering the work (Article 50(4), first subparagraph). Text published to inform the public on matters of public interest is disclosed too, unless it has undergone human review or editorial control and someone holds editorial responsibility (second subparagraph); both subparagraphs carry a law-enforcement exception. (Article 50(4)). Read: Article 50: the four limited-risk disclosures.
- Will the people affected by decisions be told AI is involved, and can they get an explanation? What a good answer includes: For an Annex III high-risk system that makes or assists decisions about natural persons, the deployer tells them it is in use (Article 26(11)); and, except for point 2 systems and subject to Article 86(2) and (3), a person affected by a decision that produces legal effects or similarly significantly affects them in a way they consider adverse to their health, safety or fundamental rights can obtain from the deployer an explanation of the AI's role and the main elements of the decision (Article 86(1)). As provider, say what your instructions for use give the deployer to do both. (Article 26(11); Article 86(1)). Read: EU AI Act deployer obligations (Article 26).
- What do your instructions for use say about the system's capabilities and limitations? What a good answer includes: For a high-risk system, instructions that carry at least the content of Article 13(3), among it the intended purpose, the level of accuracy with its metrics, robustness and cybersecurity, known or foreseeable circumstances that may lead to risks, and the human oversight measures. They must be concise, complete, correct and clear and relevant, accessible and comprehensible to deployers (Article 13(2)). (Article 13(2) and (3)). Read: What are the EU AI Act's transparency obligations toward deployers under Article 13?.

Data and GDPR

- How do you govern the data your high-risk AI is trained, validated and tested on? What a good answer includes: The practices Article 10(2) requires: design choices, data collection and origin (for personal data, the original purpose), preparation operations, assumptions, an assessment of availability, quantity and suitability, examination for possible biases and measures against them, and data gaps and how to address them. Then how the data sets meet Article 10(3) and (4): relevant, sufficiently representative, to the best extent possible free

of errors and complete, and suited, as far as the intended purpose requires, to the geographical, contextual, behavioural or functional setting of use. (Article 10(2), (3) and (4)). Read: What are the data governance requirements under the EU AI Act (Article 10)?.

- Do you process special categories of personal data to detect bias, and on what conditions? What a good answer includes: If you do, the basis is Article 4a, which Regulation (EU) 2026/1744 inserted when it deleted Article 10(5): only to the extent strictly necessary for bias detection and correction, with every condition in Article 4a(1) met in addition to the GDPR and the other data protection law it names, among them that other data, including synthetic or anonymised data, would not do. When Article 4a starts to apply is an open question, because the amending act gives it no date of its own. (Article 4a(1) and (2), as inserted by Regulation (EU) 2026/1744). Read: The data governance guide, which covers the bias-data conditions now in Article 4a(1).
- Is our data used to train your models? What a good answer includes: An answer per system from your contract with the customer and your model providers' terms. The AI Act does not answer this question, and saying so is part of a complete answer. (Not an AI Act question). Read: The AI questions supplier questionnaires ask.
- Can we use your information for our data protection impact assessment? What a good answer includes: Yes, for a high-risk system: where applicable, deployers use the information the provider gives under Article 13 to carry out their data protection impact assessment under Article 35 of Regulation (EU) 2016/679 or Article 27 of Directive (EU) 2016/680 (Article 26(9)). Point to the parts of your instructions for use that carry it. (Article 26(9)). Read: DPIA or FRIA: which one do you need?.
- Will you support our fundamental rights impact assessment? What a good answer includes: Yes, with the information Article 13 obliges you to give, which the assessment takes into account (Article 27(1), point (d)). The duty sits with deployers that are bodies governed by public law or private entities providing public services, and with deployers of systems in points 5(b) and (c) of Annex III, before deploying an Article 6(2) high-risk system outside point 2 of Annex III; under Article 27(4), as replaced, where an obligation is already met through a data protection impact assessment, the assessment can cross-refer to it or include parts of it. (Article 27(1); Article 27(4), as replaced by Regulation (EU) 2026/1744). Read: Article 27 fundamental rights impact assessment.

Oversight and monitoring

- Can a person review, override or stop the AI's output? What a good answer includes: For a high-risk system, yes by design: Article 14(4) requires the people assigned to oversight to be enabled, as appropriate and proportionate, to understand its capacities and limitations and monitor its operation, stay aware of automation bias, interpret its output correctly, decide not

to use it or disregard, override or reverse the output, and intervene in or interrupt it through a stop button or similar procedure. Name who holds the role: the deployer assigns oversight to people with the necessary competence, training and authority and the necessary support (Article 26(2)). (Article 14(4); Article 26(2)). Read: What does Article 14 of the EU AI Act require for human oversight?.

- Do your systems keep logs, and for how long? What a good answer includes: For a high-risk system, automatic recording of events over its lifetime (Article 12(1)), with the logs under your control kept for a period appropriate to the intended purpose, at least six months unless Union or national law, in particular on personal data, provides otherwise (Article 19(1)). Deployers keep the logs under their control on the same terms (Article 26(6)), and financial institutions under Union financial services law keep them within that documentation. (Article 12(1); Article 19(1); Article 26(6)). Read: What are the logging and record-keeping requirements under the EU AI Act?.
- How do you monitor the AI once it is in use? What a good answer includes: For a high-risk system, the post-market monitoring system Article 72(1) requires, proportionate to the nature of the AI technologies and the risks, which collects, documents and analyses performance data throughout the system's lifetime and, where relevant, its interaction with other AI systems (Article 72(2)), based on a plan that is part of the technical documentation (Article 72(3), as replaced). The Commission's guidance and template for that plan are due by 2 September 2027 (Article 72(3), as replaced), so name the plan you use until then. (Article 72(1) and (2); Article 72(3), as replaced by Regulation (EU) 2026/1744). Read: EU AI Act post-market monitoring (Article 72).
- How would you report a serious incident? What a good answer includes: For a high-risk system, to the market surveillance authorities of the Member States where it occurred, immediately once a causal link or its reasonable likelihood is established, and not later than 15 days after becoming aware; 10 days where a person has died, and two days for a widespread infringement or a serious and irreversible disruption of critical infrastructure (Article 73(1) to (4)). Say also how you will hear from the customer: a deployer that identifies a serious incident immediately informs the provider first (Article 26(5)). (Article 73(1) to (4); Article 26(5)). Read: AI Act serious incident reporting (Article 73).
- How do you test accuracy, robustness and cybersecurity? What a good answer includes: For a high-risk system, against the appropriate level Article 15(1) requires throughout the lifecycle, with the accuracy levels and metrics declared in the instructions for use (Article 15(3)). For security, name the measures Article 15(5) calls for where appropriate: against data poisoning, model poisoning, adversarial examples or model evasion, confidentiality attacks and model flaws. (Article 15(1), (3) and (5)). Read: What does Article 15 of the EU AI Act require for accuracy, robustness and cybersecurity?.

Documentation and evidence

- Do you have technical documentation for each high-risk system? What a good answer includes: A set per high-risk system, drawn up before it is placed on the market or put into service, kept up to date and containing at a minimum the elements of Annex IV (Article 11(1)). SMEs (including start-ups) and small mid-cap enterprises may provide those elements in a simplified manner on a form the Commission is to establish, which notified bodies accept (Article 11(1), second subparagraph, as replaced). (Article 11(1), as amended by Regulation (EU) 2026/1744; Annex IV). Read: AI Act Technical Documentation Software.
- Do you have a quality management system for your AI? What a good answer includes: For a high-risk system, a quality management system documented in written policies, procedures and instructions that covers at least the aspects Article 17(1) lists. Its implementation is proportionate to the size of your organisation, in particular for an SME or a small mid-cap enterprise, but keeps the degree of rigour and the level of protection compliance requires (Article 17(2), as replaced). (Article 17(1); Article 17(2), as replaced by Regulation (EU) 2026/1744). Read: AI Act quality management system (Article 17).
- Has the system been through conformity assessment, and can we see the declaration of conformity? What a good answer includes: For a high-risk system, which Article 43 procedure applies (internal control for points 2 to 8 of Annex III under Article 43(2), the procedure in Article 43(1) for point 1, or the procedure of the relevant act in Section A of Annex I under Article 43(3), as replaced) and the EU declaration of conformity drawn up for it under Article 47. Those provider duties (Article 16, points (f) and (g)) apply from 2 December 2027 for Annex III systems and from 2 August 2028 for Annex I systems (Article 113, third paragraph, point (c), as replaced). (Article 16, points (f) and (g); Article 43; Article 47). Read: AI Act Conformity Assessment Software.
- Is the system registered in the EU database? What a good answer includes: For a high-risk system listed in Annex III, the provider registers itself and the system before placing it on the market or putting it into service (Article 49(1)), as does a provider that concluded under Article 6(3) that an Annex III system is not high-risk (Article 49(2)); point 2 systems are registered at national level instead (Article 49(5)). Public-authority deployers, and persons acting on their behalf, register their use too (Article 49(3)). (Article 49(1), (2), (3) and (5)). Read: EU AI Act database registration (Article 49).
- Which documents will you keep for authorities, and for how long? What a good answer includes: For a high-risk system, for 10 years after it was placed on the market or put into service: the technical documentation, the quality management system documentation, the changes approved by notified bodies and their decisions and other documents where applicable, and the EU declaration of conformity (Article 18(1)). Financial institutions under

Union financial services law keep the technical documentation within that documentation (Article 18(3)). (Article 18(1) and (3)). Read: What an EU AI Act audit looks like.

STARTER

€499 a month

or €4,990 a year. Prices exclude VAT.

- ✓ Up to 20 AI systems
- ✓ Evidence and the audit log
- ✓ The Free plan is for one AI system and one user, with no end date

Get my risk classes, free

All plans

Questions

Do all thirty questions apply to every AI system? −

No. Instructions for use, data governance, human oversight, logs, monitoring, incident reporting, accuracy, technical documentation, the quality management system and conformity assessment are high-risk questions. For any other system the answer is Not applicable, with the system's classification as the reason. Inventory, roles, literacy, prohibited practices and the Article 50 disclosures reach further.

Where do these questions come from? +

Is answering all thirty the same as complying with the AI Act? +

How does Complipath help answer them? +

See what your customers will ask about your AI.

No account needed. Every answer cites the article it rests on.

Get my risk classes, free

Book a 20-min call

20 minutes with the founder.

Complipath

Get my risk classes, free

Book a 20-min call

20 minutes with the founder.

Product

- Get my risk classes, free
- Pricing
- All features
- What Complipath does not do
- Risk classification
- AI inventory
- Annex IV documentation
- Sign in

Resources

- All guides
- Templates
- Sources
- Deadlines
- Why now
- What the check can and cannot decide
- What we cover of the Act
- Free readiness review
- For small companies
- For SaaS companies
- For AI providers
- How to choose
- Compare tools
- Glossary
- FAQ
- Changelog

Why not just ask a model

Do we still need a lawyer?

LinkedIn

Company

What Complipath is

Compare

What to ask us before you buy

What you can check before you trust us

Security

Security pack

About

Press

Partners

hello@complipath.io

Legal

Privacy

Terms

DPA

SLA

Subprocessors

Imprint

PAGE complipath.io/ai-act-customer-questions **BUILT** 2026-11-03 **CLASSIFIER** 2026-10-23.1

LAW VERSION 2024/1689 + 2026/1744 + Commission guidelines

Complipath is EU AI Act compliance software for companies with 20–100 employees that use or build AI. A rule engine classifies each AI system and cites the article; AI only drafts from your own documents, and a person confirms.

We measure page views, and clicks on links into the app or to booking a call, with Vercel Web Analytics. It uses no third-party cookies. Visitors are identified by a hash derived from the incoming request, which is discarded after 24 hours, and no identifier is stored that could follow a visitor to another site. What is collected: the time of the visit, the URL, the referring page, filtered query parameters, city-level location, operating system, browser and device type.

Complipath AB · Org. no. 559609-5835 · Holländareplatsen 14, 415 14 Göteborg, Sweden · hello@complipath.io

Complipath (complipath.io), EU AI Act compliance software from Gothenburg, Sweden.

Complipath is built by Yobel Tzegai in Gothenburg, Sweden.

Complipath provides legal information, not legal advice. Every guide cites its source on EUR-Lex — Regulation (EU) 2024/1689, and Regulation (EU) 2026/1744 where that has amended it; where the law is still settling, the guide says so.