

GRECTA

One standard, every regime

AI ACT TRANSPARENCY & READINESS PACK

Ready-to-use Article 50 transparency notices, Article 26 deployer notices, an Article 4 AI-literacy starter, and a 15-question buyer due-diligence bank. Built for providers and deployers.

Document	Marketing pack — transparency notices and readiness starters
Edition	October 2026
Publisher	GRECTA
Built for	Providers and deployers of AI systems placing them on or using them in the EU market
Covers	AI Act Article 50 (transparency), Article 26 (deployer obligations), Article 4 (AI literacy); plus a buyer due-diligence question bank
Scope limit	This is a transparency and readiness starter. It is NOT a substitute for Annex IV technical documentation, a conformity assessment, or a fundamental-rights impact assessment
Reviewed against	The consolidated text of Regulation (EU) 2024/1689 (as amended by the Digital Omnibus, Regulation (EU) 2026/1744) and Commission guidance current to October 2026
GRECTA	One standard, every regime

This is a free marketing resource published by GRECTA to help providers and deployers get started on the AI Act's transparency and readiness obligations. It is not legal advice, is not an official document, and is not approved by the European Commission, the AI Office or any authority. The notices are starting templates to be adapted to your systems and reviewed by a qualified person. This pack is not a substitute for the Annex IV technical documentation, conformity assessment or impact assessment that the Act separately requires for high-risk systems. It was reviewed against the consolidated text of Regulation (EU) 2024/1689, as amended by the Digital Omnibus (Regulation (EU) 2026/1744), and Commission guidance current to October 2026; the AI Act moves quickly, so verify the current position before you rely on anything here.

About this pack

- (1) Most organisations meet the AI Act first not through its heavy high-risk machinery but through three lighter obligations that arrive sooner and touch almost everyone: the duty to be transparent when people deal with AI (Article 50), the duties that fall on an organisation when it deploys someone else's AI system (Article 26), and the duty to make sure the people operating AI on your behalf actually understand it (Article 4, AI literacy). On top of those, any organisation that sells software increasingly meets a fourth thing — a customer's due-diligence questionnaire asking how it handles AI. This pack gives you a ready-to-use starting point for all four.
- (2) It is deliberately a starter, not a compliance system. The notices here are drafted to be adapted to your own systems and dropped into your product, your contracts and your onboarding; the literacy starter gets a programme off the ground; the question bank lets you answer — or send — a buyer questionnaire with confidence. What it is not is the deeper evidence the Act requires for high-risk systems: the Annex IV technical documentation, the conformity assessment, the risk-management file and the fundamental-rights impact assessment are separate exercises, and this pack is explicit about where its edge is. GRECTA exists to carry an organisation from these starters through to that deeper, cross-regime governance — one classification of what your product is and does, mapped once and reused across every regime it touches.

WHERE THIS PACK STOPS

A starter, not a technical file.

This pack covers the transparency, deployer and literacy obligations and buyer due diligence. It does not produce the Annex IV technical documentation, the Article 9 risk-management file, the Article 11 documentation, the conformity assessment, or the Article 27 fundamental-rights impact assessment that high-risk systems separately require. Treat a completed notice here as a public-facing or contractual artefact, not as evidence of high-risk conformity.

A note on the Digital Omnibus timeline.

The Digital Omnibus (Regulation (EU) 2026/1744, in force 27 July 2026) deferred the high-risk application dates to December 2027 and August 2028, but it did NOT defer the obligations in this pack: Article 50 transparency applies from 2 August 2026, and the Article 26 deployer and Article 4 literacy duties are unchanged. A delay to high-risk is not a reason to stand down on these — for most organisations, these are the duties that bite first.

Section 1

Article 50 transparency notices

1. What Article 50 requires

- (3) Article 50 sets transparency obligations that apply regardless of risk tier, and they fall on providers and deployers in different ways. A provider of an AI system intended to interact directly with people must ensure the system is designed so that people are informed they are interacting with AI, unless that is obvious to a reasonably well-informed person. A provider of a generative AI system must ensure its synthetic audio, image, video or text output is marked in a machine-readable format as artificially generated or manipulated. A deployer of an emotion-recognition or biometric-categorisation system must inform the people exposed to it. And a deployer that generates or manipulates deep-fake image, audio or video content, or text published to inform the public on matters of public interest, must disclose that the content is artificially generated or manipulated. The information must be given clearly and at the latest at the first interaction or exposure.
- (4) The notices below are ready-to-use starting text for each of these situations. Adapt the wording to your product and brand voice, place each where the law requires (at first interaction, at exposure, or with the content), and keep a record of where you have deployed it. The machine-readable marking duty for generative output is a technical measure as well as a notice — the notice tells the human, the marking tells the machine — so the generative notice below is paired with a reminder about the marking itself.

RELEVANT LEGAL TEXT

Article 50 AI Act — Transparency (abridged).

Providers shall ensure that AI systems intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system, unless this is obvious from the point of view of a reasonably well-informed, observant and circumspect natural person. Providers of AI systems generating synthetic audio, image, video or text content shall ensure the outputs are marked in a machine-readable format and detectable as artificially generated or manipulated. Deployers of emotion recognition or biometric categorisation systems shall inform the natural persons exposed. Deployers generating or manipulating deep-fake content shall disclose that the content has been artificially generated or manipulated; where the content is text published to inform the public on matters of public interest, its artificial generation or manipulation shall be disclosed. Information shall be provided in a clear and distinguishable manner at the latest at the time of the first interaction or exposure.

2. Notice A — AI interaction disclosure (provider; chatbots, assistants, agents)

For a system that interacts directly with people. Show at or before first interaction. Omit only if it is genuinely obvious the user is dealing with AI.

Short form (in-product)	“You’re chatting with an AI assistant, not a human.”
Standard form	“This is an AI system. You are interacting with automated software, not a person. It can make mistakes — check anything important.”
Extended form (first-run / help)	“[Product] is an artificial intelligence system. Your messages are processed by automated software to generate responses. It is not a human and may produce inaccurate or incomplete information; do not rely on it for [legal / medical / financial] decisions without professional advice. [Link: how this works / data use].”

3. Notice B — AI-generated content marking (provider; generative systems)

Pair the human-readable label below with the machine-readable marking the Act requires in the output itself (e.g. metadata / watermarking). The label is not a substitute for the marking.

Visible label	“Generated with AI.” / “AI-generated. May not depict real events.”
With provenance	“This [image / audio / video / text] was generated or edited using AI by [Product]. A machine-readable marker identifying it as AI-generated is embedded in the file.”
Machine-readable reminder	Ensure the output carries a machine-readable and detectable marker of artificial generation/manipulation (e.g. C2PA/content credentials, metadata, or watermarking), interoperable where feasible. The visible label tells people; the marker satisfies the technical duty.

4. Notice C — emotion recognition / biometric categorisation (deployer)

Inform the people exposed, before or at exposure. (Note: some emotion-recognition and biometric uses are prohibited or high-risk — confirm your use is permitted before relying on a notice.)

Notice	“[Organisation] uses an AI system that [recognises emotional states / categorises people using biometric data] in [context]. You are being informed because you are within its range. [What is processed, why, and for how long]. [How to object / alternative route, where available].”
---------------	--

5. Notice D — deep-fake / AI-manipulated content disclosure (deployer)

Disclose with the content. For text published to inform the public on matters of public interest, disclose unless the content has undergone human review and a person holds editorial responsibility.

Image / audio / video	“This content has been artificially generated or manipulated using AI.” — placed visibly on or immediately with the content.
Public-interest text	“Parts of this text were generated or altered using AI.” — where applicable and not subject to the human-editorial-responsibility exception.
Artistic / creative exception	Where the content is evidently part of a creative, satirical, artistic or fictional work, the disclosure may be made in a manner that does not hamper the display or enjoyment of the work.

6. Deployment record

Keep a short record of where each notice is live — useful if an authority or a customer asks.

System	Notice used (A/B/C/D)	Where it appears	Live since

Section 2

Article 26 deployer notices

1. What Article 26 requires of deployers

- (5) When an organisation uses a high-risk AI system that someone else provides, it is a deployer, and Article 26 places a distinct set of duties on it — separate from the provider’s. A deployer must use the system in accordance with its instructions for use; assign human oversight to people with the necessary competence, training and authority; ensure that input data within its control is relevant and sufficiently representative; monitor the system’s operation and suspend use and inform the provider where it presents a risk or a serious incident occurs; keep the automatically generated logs it controls; and, in defined cases, inform affected people. Where the deployer is a public body or provides public services, or deploys certain systems, it may also owe a fundamental-rights impact assessment under Article 27 and registration duties. The notices and records in this section help a deployer put the visible, people-facing and provider-facing parts of Article 26 in place.
- (6) These are starters for the deployer’s own notices and internal records — the things a deployer issues or keeps, as distinct from the provider’s technical file. They are written so a deployer can inform its workforce where an AI system is used in the workplace, inform individuals where a high-risk system makes or supports decisions about them, and keep a simple internal record that oversight, monitoring and logging are actually assigned and happening. They do not replace the Article 27 impact assessment, which is a separate exercise where it applies.

RELEVANT LEGAL TEXT

Article 26 AI Act — Obligations of deployers (abridged).

Deployers of high-risk AI systems shall take appropriate technical and organisational measures to ensure they use such systems in accordance with the instructions for use. They shall assign human oversight to natural persons who have the necessary competence, training and authority, as well as the necessary support. To the extent the deployer exercises control over the input data, it shall ensure that input data is relevant and sufficiently representative in view of the intended purpose. Deployers shall monitor the operation of the high-risk AI system on the basis of the instructions for use and, where relevant, inform providers; where they identify a risk or a serious incident they shall suspend use and inform the provider. They shall keep the logs automatically generated by the system that are under their control. Before putting into service or using a high-risk AI system at the workplace, deployers who are employers shall inform workers’ representatives and the affected workers that they will be subject to its use. Where a high-risk system makes or assists in decisions related to natural persons, deployers shall inform those persons that they are subject to its use.

2. Notice E — workplace deployment notice (to workers)

Give before putting a high-risk AI system into use at work; inform workers’ representatives and affected workers.

Notice

“[Organisation] will begin using an AI system, [name], in [function / team] from [date]. It is used to [purpose]. It [supports / informs] decisions about [what]; final decisions rest with [role]. Human oversight is assigned to [role]. If you have questions or wish to raise a concern, contact [contact]. [Where applicable: this has been shared with workers’ representatives on (date).]”

3. Notice F — affected-person notice (high-risk decision about an individual)

Give where a high-risk system makes or assists a decision about a person (e.g. recruitment, credit, access to services).

Notice

“An AI system is used to help [assess / decide] [what] in your case. It [does /

	does not] make the final decision — [role] is responsible for the outcome. You have the right to [an explanation / human review / to contact us] at [contact]. [Link: how the system is used and your rights.]”
Explanation route (Art. 86)	Be ready to give a clear, meaningful explanation of the role the system played in a decision that produces legal or similarly significant effects, where a person asks.

4. Deployer oversight & monitoring record

A simple internal record that the Article 26 duties are actually assigned and happening. Not the provider's technical file.

System and provider	<i>The high-risk system and who provides it.</i>
Used per instructions for use?	<i>Confirmation the system is used within its intended purpose and instructions.</i>
Human oversight assigned to	<i>The named role(s), and confirmation of competence, training and authority to override or stop.</i>
Input data responsibility	<i>Where the deployer controls input data, how it is kept relevant and representative.</i>
Monitoring	<i>How operation is monitored, and the route to suspend use and inform the provider on a risk or serious incident.</i>
Logs	<i>Where the automatically generated logs under the deployer's control are kept, and for how long.</i>
Article 27 FRIA needed?	<i>Whether a fundamental-rights impact assessment applies (public body / public-service provider / listed use). Separate exercise — not in this pack.</i>
Registration	<i>Whether deployer registration applies, and its status.</i>

Section 3

Article 4 AI-literacy starter

1. What Article 4 requires

- (7) Article 4 is short and easy to underestimate: providers and deployers must take measures to ensure, to their best extent, a sufficient level of AI literacy among their staff and other people operating and using AI systems on their behalf — taking into account those people’s technical knowledge, experience and training, the context the systems are used in, and the people the systems are used on. It has applied since February 2025, it has no risk-tier threshold, and it reaches any organisation that puts AI in the hands of its people. There is no prescribed curriculum; what is required is a reasonable, proportionate effort that you can show you made. This starter gives you that: a minimal programme and a record that it happened.
- (8) The point of literacy is not certification; it is that the people who operate AI can recognise what it is good and bad at, spot when it is going wrong, understand the specific risks of the systems your organisation actually uses, and know what the law expects of them. A proportionate programme for a small team is modest — a short session, role-specific pointers, and a note of who attended. The starter below scales from that up to a structured rollout, and pairs with the deployer oversight duties in Section 2, because the people exercising oversight are exactly the people who most need the literacy.

RELEVANT LEGAL TEXT

Article 4 AI Act — AI literacy.

Providers and deployers of AI systems shall take measures to ensure, to their best extent, a sufficient level of AI literacy of their staff and other persons dealing with the operation and use of AI systems on their behalf, taking into account their technical knowledge, experience, education and training and the context the AI systems are to be used in, and considering the persons or groups of persons on whom the AI systems are to be used.

2. A minimal programme that satisfies the duty

- (9) The following is a proportionate baseline. Scale the depth to the roles involved — heavier for those exercising oversight of high-risk systems, lighter for occasional users of low-risk tools.

Element	What to cover
What AI is, and is not	Plain explanation of the AI systems your organisation uses; what they do; their limits; that they can be confidently wrong.
The specific risks here	The risks of your actual systems — bias, error, over-reliance, data leakage, hallucination — with concrete examples from your context.
Your rules	What staff may and may not use AI for; approval routes; what must never go into a tool; who to ask.
The law, briefly	That the AI Act applies; what transparency, oversight and (where relevant) high-risk duties mean for their work.
Oversight & escalation	For those overseeing AI: how to monitor it, when and how to override or stop it, and how to report a problem or serious incident.
Refresh	A light annual refresh and an update when a significant new system or risk appears.

3. Literacy record

The evidence that you made the effort. This is what you show if asked.

Programme owner	<i>Who is responsible for AI literacy.</i>
Audience and tailoring	<i>Which roles were trained, and how depth was matched to their work and the systems they touch.</i>
Delivered	<i>What was delivered, how (session / module / briefing), and when.</i>
Attendance	<i>Who attended / completed, kept as a list or register.</i>
Oversight roles	<i>Confirmation that those exercising human oversight received the deeper content.</i>
Next refresh	<i>DD.MM.YYYY, and the trigger for an earlier update.</i>

4. Half-page briefing you can hand out

A ready starter you can adapt into a one-page staff handout.

“We use AI tools at [Organisation]. They can help, but they can also be wrong in confident-sounding ways. Before you rely on AI output, check it. Never put [confidential / personal / client] information into a tool unless it is on the approved list. If an AI system is making or supporting a decision about a person, a named human is responsible for that decision — not the tool. If something looks wrong, stop and tell [contact]. The law (the EU AI Act) expects everyone operating AI here to understand these basics — that is why this matters.”

Section 4

Buyer AI due-diligence question bank (15)

1. How to use the question bank

- (10) When one organisation buys or integrates another's AI, the buyer runs due diligence — and increasingly sends a questionnaire. This bank of fifteen questions works in both directions. If you are a buyer or a deployer procuring AI, send it to your vendor and judge the answers. If you are a provider, answer it in advance so you can respond to a customer's questionnaire in hours rather than weeks — the thing that so often stalls an enterprise deal. The questions are mapped to the AI Act article they probe, so an answer that is thin on a given question tells you precisely which obligation to look at.
- (11) Good answers are specific and evidenced, not reassuring adjectives. “Yes, we take bias seriously” is not an answer; “we test for X across Y groups, last run on Z, results held in our risk file” is. The right-hand column notes what a strong answer looks like, so the bank doubles as a scoring guide.

BOTH DIRECTIONS

Buyer / deployer: send it and score the answers.

Use the “what a strong answer shows” column as your scoring guide. A vague answer is a finding, not a pass.

Provider: answer it in advance.

Keep your answers current and evidenced. A pre-answered bank is the fastest way to clear a customer's security-and-AI questionnaire without stalling the deal.

2. The fifteen questions

#	Question	What a strong answer shows	Art.
1	What exactly does the AI system do, and what is its intended purpose?	A precise purpose, not marketing; the boundary of intended use stated.	3,13
2	How is the system classified under the AI Act — prohibited, high-risk, transparency-only, GPAI or minimal — and on what basis?	A reasoned classification citing Annex III / criteria, not a bare label.	6,50
3	Are you the provider of this system, or a deployer/integrator of someone else's?	A clear role; if they integrate a third-party model, who answers for what.	3,25
4	If high-risk, is there Annex IV technical documentation and a completed conformity assessment?	The documentation exists and can be evidenced; route and any notified body named.	11,43
5	How do you meet the Article 50 transparency duties (AI interaction, AI-content marking)?	Concrete notices and, for generative output, machine-readable marking in place.	50
6	How is human oversight designed and exercised?	Specific oversight measures and real authority to override or stop — not “a human is in the loop”.	14
7	What is your approach to risk management across the lifecycle?	An actual iterative process with a risk file, not a one-off assessment.	9
8	How do you govern training, validation and test data — provenance and bias?	Provenance documented; bias examined and mitigated; representativeness assessed.	10
9	What are the system's known limitations, error rates and failure modes?	Honest, quantified limits and the conditions under which it degrades.	13,15

10	How accurate, robust and secure is it, and how is that measured?	Named metrics, test results, and security measures against adversarial misuse.	15
11	How do you log system activity, and what logs can we obtain as a deployer?	Automatic logging; what the deployer receives for its own Article 26 duties.	12,26
12	What is your serious-incident process, and will you notify us?	A defined Article 73 process and a contractual commitment to notify in time.	73
13	What do you provide deployers to meet their own obligations (instructions, information)?	Clear instructions for use and the information a deployer needs to comply.	13,26
14	How do you ensure AI literacy among those operating the system?	A real, proportionate programme with a record — not an assertion.	4
15	How do you keep compliance current as the system and the law change?	Version control, monitoring of legal change, and an update commitment to customers.	72

Question 4 is the sharpest filter: if a system is high-risk and there is no Annex IV technical documentation and conformity assessment, nothing else in the list compensates. This pack does not produce that documentation — it is a separate exercise.

From starter to standard

- (12) These four starters cover the obligations most organisations meet first. The harder part is everything underneath and alongside them — the Annex IV technical documentation, the risk and impact files, and the fact that the same product usually sits under several regimes at once (the AI Act, GDPR, DORA, NIS2, ISO 42001). Handling each law as its own project is how compliance cost runs away. GRECTA takes the opposite approach: one classification of what your product is and does, mapped once and reused across every regime it touches. One standard, every regime.

START WITH A FREE SCOPING CALL

Begin a free pilot — see where your products actually stand.

In a free scoping call we classify one of your AI products and show you, across the AI Act and the other regimes it touches, exactly which obligations reach you and where the gaps are — including the Annex IV and conformity work this starter pack deliberately leaves out. No cost, no obligation.

Two ways to start.

Book directly: <https://calendly.com/yuliia-grecta/30min> — or send us a message through the contact page at grecta.com, or reply to wherever you received this pack. Tell us what you build; we will bring the map of what governs it.

GRECTA

One standard, every regime. This pack is a free marketing resource, not legal advice, and is not a substitute for Annex IV technical documentation or a conformity assessment. Reviewed against the consolidated text of Regulation (EU) 2024/1689 (as amended by Regulation (EU) 2026/1744) and Commission guidance current to October 2026; verify the current position before relying on it.