

>_ EU AI ACT · ARTICLE 4 · 2026 EDITION

The AI Literacy Starter Kit

Article 4 has applied to every provider and deployer since 2 February 2025, whatever the risk tier. In July 2026 it was rewritten. You no longer have to reach a "sufficient level" of AI literacy, but you still have to take measures, and be able to show them. This kit covers what that means and how to get it done in ninety days, without buying a course nobody finishes.

>_ WHAT CHANGED IN JULY 2026

Regulation (EU) 2026/1744 replaced Article 4. The duty moved from ensuring a level of literacy to taking measures that support it. The wording got softer. The need for a record did not.

01 · THE OBLIGATION

What Article 4 asks for now

Regulation (EU) 2026/1744, the "Digital Omnibus on AI", replaced Article 4 in full. It entered into force on 27 July 2026 and the new text applies from that day.

ARTICLE 4(1), AS AMENDED · OUR PARAPHRASE

Providers and deployers of AI systems must take measures to support the development of AI literacy among their staff and anyone else who operates or uses AI systems on their behalf. The measures should reflect those people's technical knowledge, experience, education and training, the context the systems are used in, and the people the systems are used on. The duty does not require anyone to guarantee a specific level of AI literacy in any individual.

A paraphrase, not the official text. The short phrases quoted below match the amended wording as published in consolidated versions of the Act. Check them against the Official Journal text on EUR-Lex before you quote them in a policy.

Four phrases that decide what you have to build

THE PHRASE	WHAT IT MEANS IN PRACTICE
"take measures to support the development of AI literacy"	A duty of effort, not of result. Nobody will test whether your staff reached a level; they will ask what you did. Measures you cannot show look exactly like no measures, so record them.
"does not require ... to guarantee any specific level"	One person's bad call does not put you in breach if the programme was reasonable. Having no programme does.
"other persons ... on their behalf"	Not just employees. Contractors, agency staff and outsourced teams who run your systems are in scope. Unchanged by the rewrite.
"the context the AI systems are to be used in"	Still the clause that matters most. One generic company-wide course does not fit a duty tied to context. A recruiter and a copywriter need different things.

UNTIL 26 JULY 2026

Providers and deployers had to "ensure, to their best extent, a sufficient level of AI literacy". That read like a target you could miss.

FROM 27 JULY 2026

Take measures that support AI literacy, sized to context. A new paragraph asks the Commission and Member States to support these efforts, particularly for SMEs, including practical examples.

What it is not

- **Not optional because it got softer.** It still binds every provider and deployer, at every risk tier. If your people use a chatbot at work, it applies.
- **Not a certification scheme.** There is no accredited "AI literacy" badge you can buy your way to.
- **Not a one-off.** Systems and models change, and your record should show refreshes.
- **Not a policy people click through.** A policy states rules; literacy is the ability to apply them to a situation nobody wrote down.

02 · SCOPE

Who counts, and who is usually missed

Start by writing the list down. Most organisations find they were thinking about a third of the people who are actually in scope.

GROUP	WHY THEY'RE IN SCOPE
Direct users	Anyone who prompts a model, in any tool, including the ones procurement never saw.
Human overseers	People named as reviewers on a high-risk system. Article 26(2) separately requires them to have the competence, training and authority for the role. For Annex III systems, that applies from 2 December 2027.
Contractors & agency staff	"Other persons ... on their behalf". Explicitly covered, routinely forgotten.
Builders	Anyone configuring agents, prompts or connectors. A substantial modification to a high-risk system can make you its provider under Article 25.
Procurement & vendor managers	They decide which systems come in, and what documentation arrives with them.
Executives & board	Accountable for the programme, rarely in the tool. They need a different twenty minutes from everyone else.

The uncomfortable first step

You cannot train people on systems you haven't listed. Before any curriculum, build a register of the AI systems actually in use, including personal subscriptions and free tiers. Expect the real number to be well above the number IT knows about.

That gap is not a failure of policy. It is what happens when the approved route is slower than the unapproved one, and it tells you where to start: make the approved route the easy one, then teach people how to use it well.

>_ A USEFUL TEST

Ask five people in different departments to name every AI tool they used last week. If the answers surprise you, your literacy programme has a scoping problem before it has a content problem.

03 · CURRICULUM

Six roles, six different needs

Sized to context means the recruiter and the copywriter get different material. Here is a defensible starting split. Adapt the roles to your org chart, not the other way round.

ROLE	WHAT THEY NEED TO KNOW	TIME
People & hiring	Why CV screening and candidate ranking are high-risk under Annex III, point 4(a), with deployer duties applying from 2 December 2027; what human oversight means in a hiring decision; candidate data handling; the duty to inform workers' representatives (Art. 26(7)).	45 min
Marketing & brand	The Article 50 transparency duties, which apply from 2 August 2026; when a deepfake must be disclosed; how generated images, audio and video get marked; backing up claims made in generated copy.	30 min
Legal & compliance	Classification and the Annex III list; the revised timeline in Regulation (EU) 2026/1744; when a fundamental rights impact assessment is needed (Art. 27); record-keeping and retention; serious-incident reporting (Art. 73).	90 min
Engineering	When configuration becomes substantial modification (Art. 25); logging; model version pinning and change management; input data quality (Art. 26(4)).	60 min
Finance & operations	Approval steps on actions with consequences; automation bias; input data relevance; retention and legal hold.	30 min
Executives	Where liability actually sits; the penalty tiers in Art. 99 and how they scale for smaller companies; what to ask for in a quarterly review; how to read the register.	20 min

The universal twenty minutes

Everyone, whatever their role, should be able to answer four questions:

- What are the prohibited practices, and why does no business case override them? (Two were added in 2026; see the dates on page 6.)
- What counts as personal data in a prompt?
- How do you tell a confident wrong answer from a right one?
- Who do you ask when you aren't sure?

If a module doesn't help people answer one of those four, cut it.

04 · ROLLOUT

Thirty, sixty, ninety

The goal of the first ninety days is not a trained workforce. It is a record you can defend, a closed shadow-AI gap, and a habit that outlasts the launch.

1-30

DAYS

SEE IT

Build the register and the baseline

- List every AI system in actual use, including personal subscriptions. Give each one an owner.
- Name one accountable owner for the literacy programme itself. One person, not a committee.
- Publish the AI use policy where the work happens, not on an intranet page.
- Run the universal twenty minutes for everyone. Record completion: this is your baseline evidence.
- Find out who is currently named as a human overseer on anything high-risk. Expect gaps.

31-60

DAYS

AIM IT

Role paths and the first gates

- Ship the role-specific modules for the two or three roles carrying the most exposure.
- Make training a condition of oversight: no training record, no approver role. Annex III oversight duties apply from 2 December 2027, and a gate built now is cheaper than one retrofitted then.
- Start a scenario bank from your own blocked prompts and near-misses, anonymised.
- Brief executives separately. Twenty minutes, no jargon, focused on where liability sits.
- Bring contractors and agency staff into the same path as employees.

61-90

DAYS

MAKE IT
STICK

Teaching in context, and refresh cycles

- Move teaching to where prompts are written: short, specific guidance at the moment of the decision.
- Set refresh intervals: twelve months for general literacy, shorter for high-risk oversight.
- Publish adoption and quality metrics by team, never by individual.
- Add a literacy line to the quarterly governance review, next to the register and the incident log.
- Run a tabletop: an auditor asks for evidence of AI literacy. Time how long it takes to produce.

Why the paperwork matters more after the rewrite

The amended Article 4 judges the measures you took, not a level anyone reached, so the record is the whole case. Every step above should leave something behind: a list, a date, a name, a completion log.

05 · TEACHING MATERIAL

Twelve scenarios worth more than a slide deck

Scenario questions beat multiple-choice recall, because the failure you care about is judgement under time pressure. Use these as written, or as a template for ones drawn from your own logs.

01 The pasted spreadsheet PERSONAL DATA

A colleague pastes a customer export into a prompt to "find patterns". What has just happened, and what should they have done instead?

02 The confident citation AUTOMATION BIAS

An answer cites a regulation section that reads perfectly and does not exist. Which three checks would have caught it?

03 The shortlist ANNEX III, 4(A)

A manager asks the model to "rank these twelve applicants". Where does this become high-risk, and what changes if it only summarises?

04 The sentiment dashboard ART. 5(1)(F)

HR proposes scoring team morale by inferring emotions from how people write. Why will no risk assessment rescue this?

05 The unlabelled campaign ART. 50

A realistic generated image of a real person ships in a campaign with no disclosure. What was required, and at which step in the workflow?

06 The helpful agent APPROVAL STEPS

An agent is given write access to the CRM to "save time". What should have gone to a human first?

The dates these scenarios lean on

2 Feb 2025 Article 4 (AI literacy) and the original Article 5 prohibitions apply.

27 Jul 2026 Regulation (EU) 2026/1744 in force. The rewritten Article 4 applies.

2 Aug 2026 Article 50 transparency duties apply. National supervision and enforcement of Article 4 starts in August 2026.

2 Dec 2026 Two new prohibitions apply: AI that generates non-consensual intimate imagery, and AI that generates child sexual abuse material. Generative systems already on the market before 2 Aug 2026 must meet the Article 50 marking duty.

2 Dec 2027 Annex III high-risk rules apply (hiring, worker management, education, credit, essential services and more), with the deployer duties in Articles 26, 27 and 86.

2 Aug 2028 High-risk rules for AI in products covered by Annex I (machinery, medical devices, toys and similar).

05 · TEACHING MATERIAL, CONTINUED

Six more, aimed at the builders

The first six cover everyday use. These six cover the failures that only happen once someone has enough access to cause them.

07 The silent upgrade CHANGE MANAGEMENT

The provider moves a model alias to a new version overnight. What in your assessment just went out of date?

08 The rubber stamp ART. 14(4)(B)

A reviewer has approved 340 runs with a median review time of four seconds. Is that oversight?

09 The free tier SHADOW AI

A team adopts a free tool on a personal card because the approved one was slower. What is the actual root cause?

10 The fine-tune ART. 25

Engineering fine-tunes an open model for a hiring workflow and puts it into service under its own name. Which role is the company now in?

11 The explanation request ART. 86

A rejected applicant asks why. Once the Annex III rules apply in December 2027, what do you need to have kept to answer them?

12 The contractor SCOPE

An agency operates your support bot. Whose literacy obligation is that, and who keeps the record?

> SCORING TIP

Mark the reasoning, not the answer. Someone who reaches the wrong conclusion for defensible reasons is safer than someone who guesses right.

Write your own

The best scenarios come from your own near-misses. Strip out names and customer details, keep the moment of decision, and ask what the person should have done next.

06 · EVIDENCE

What to keep on file

A duty to take measures is only as good as your record of the measures. If an authority, a customer or a works council asks tomorrow, this is the folder.

- ☐ **The AI system register**
Every system in use, with an owner, a purpose, and a classification decision with its reasoning.
- ☐ **Curriculum and rationale**
Which roles get which material, and why that split fits their context. Write the reasoning down once.
- ☐ **Completion records per person**
Who completed what, when, and when it expires. Include contractors and agency staff.
- ☐ **Competence evidence for overseers**
For every named human overseer on a high-risk system: their training, and the authority they hold to stop a run.
- ☐ **Policy acceptance and re-attestation**
Acceptance on joining, and again each time the policy materially changes.
- ☐ **Gap analysis and remediation**
Which teams lagged, what you did about it, and whether it worked. A gap found and closed is stronger evidence than a flawless first pass.
- ☐ **Refresh cycle**
The stated interval, plus proof that the second round actually happened.
- ☐ **Board or governance minutes**
Evidence that literacy was reviewed by people with the authority to fund it.

One test worth running

Pick a random employee and a random AI system. How long does it take to produce their training record, the system's classification and the name of its human overseer? If the answer is measured in days, the problem is not your curriculum. It is that the evidence is not a by-product of the work.

If you use StickyPrompts. Admins can allow or block model providers per team, and the Privacy Log records guardrail triggers and policy changes, with an export. That gives you part of the record above. It does not train anyone, and no tool makes an organisation compliant on its own.

Not legal advice. A plain-language orientation to Article 4 of Regulation (EU) 2024/1689 (the AI Act) as amended by Regulation (EU) 2026/1744 (published in the Official Journal on 24 July 2026, in force 27 July 2026), and to related provisions. Last reviewed 28 September 2026. Article text is paraphrased unless shown in quotation marks. Check current requirements against the consolidated text on EUR-Lex and with your own counsel before relying on anything here.

© 2026 StickyPrompts · LogiNet Enterprise Software Technologies Kft. Share it freely, no sign-up asked for. · stickyprompts.com/eu-ai-act