WHITEPAPER · NOVEMBER 2026

The EU AI Act in Practice: what Irish organisations need to do now

AI literacy, human oversight and the new deadlines after the 2026 Digital Omnibus. For public bodies, SMEs and the people who buy AI.

Aug 2026

AI literacy is enforceable and the AI Office of Ireland is open

Article 4; Regulation of AI Act 2026

Dec 2027

High-risk rules apply to Annex III systems, including public services and hiring

Digital Omnibus, Reg. (EU) 2026/1744

€35m

Top fine for prohibited AI practices, or 7% of worldwide turnover if higher

Article 99

> *Most organisations will never build an AI system. Almost all of them already use one. The Act's everyday duties fall on the user, and the first of them is already in force.*

Executive summary

The EU AI Act is no longer a future regulation. Two of its obligations already apply to almost every Irish organisation that uses AI, and since August 2026 Ireland has a statutory AI Office and a set of regulators ready to supervise them.

2026 changed the picture in two ways. First, the EU's **Digital Omnibus on AI** (Regulation (EU) 2026/1744, in force since 27 July 2026) moved the start date for most high-risk obligations from August 2026 to **2 December 2027**, softened the wording of the AI literacy duty, added a new prohibition on AI that generates non-consensual intimate images and child sexual abuse material, and extended SME relief to small mid-caps. Second, Ireland passed the **Regulation of Artificial Intelligence Act 2026**, which established the **AI Office of Ireland** on 2 August 2026 as the single point of contact, coordinating a distributed group of existing regulators.

The delay to the high-risk rules is real, but it is not a pause. Three things apply now:

- **Prohibited practices** (since 2 February 2025), such as social scoring and emotion recognition in the workplace.
- **AI literacy** (since 2 February 2025; supervised by national authorities from 2 August 2026): organisations must take measures to build the AI skills of the staff and contractors who use AI on their behalf.
- **Transparency** (since 2 August 2026): people must be told when they are dealing with a chatbot, and synthetic content must be marked.

The organisations that will find December 2027 easiest are the ones that use the next twelve months to do four things: list every AI system they use, decide which ones are high-risk, train their people in proportion to that risk, and write the Act's requirements into the contracts with the companies that supply their AI.

Who this paper is for

Leaders, managers and procurement teams in Irish public bodies and in small and medium-sized organisations that use, or are about to buy, AI. It sets out eight key factors, a role-based approach to AI literacy, what public bodies must add on top, the questions to ask suppliers, a 90-day plan and a 15-point checklist.

What it is not

Legal advice. The Act is long, the guidance is still arriving and the details depend on your systems. Use this paper to organise the work and frame the questions; take legal advice on specific decisions.

The timeline after the Digital Omnibus

The AI Act (Regulation (EU) 2024/1689) entered into force on 1 August 2024 and applies in stages. The Digital Omnibus changed several dates. This is where things stand in November 2026.

DATE	WHAT APPLIES
2 Feb 2025	**Prohibited practices** (Article 5) and **AI literacy** (Article 4) apply. IN FORCE
2 Aug 2025	Rules for **general-purpose AI models**, governance and penalties apply. Member States designate their authorities. IN FORCE
27 Jul 2026	**Digital Omnibus** enters into force: amended AI literacy wording, new dates for high-risk rules, SME and small mid-cap relief, bias-data provisions. IN FORCE
2 Aug 2026	**Transparency** duties (Article 50) apply. National authorities begin supervising AI literacy. The **AI Office of Ireland** is established under the Regulation of Artificial Intelligence Act 2026. IN FORCE
2 Dec 2026	New prohibition on AI generating non-consensual intimate images and child sexual abuse material. Generative systems already on the market must mark their output (Article 50(2)). WEEKS AWAY
Aug 2027	Each Member State must have at least one AI regulatory sandbox operating. NEXT YEAR
2 Dec 2027	**High-risk rules for Annex III systems**: employment, education, access to essential public and private services, critical infrastructure, law enforcement, migration and justice. Includes the deployer duties in Article 26 and the fundamental rights impact assessment in Article 27. PLAN NOW
2 Aug 2028	High-risk rules for AI built into products covered by EU product law (Annex I), such as medical devices and machinery components. LATER
2 Aug 2030	Deadline for **public authorities** to bring high-risk systems already in use before the rules applied into line with the Act (Article 111(2)). LATER

A note on official pages

Some pages written before July 2026 still show the original dates (high-risk rules from August 2026 and 2027). If training or supplier material still quotes August 2026 for high-risk systems, it is out of date.

Start with two questions

Almost every obligation in the Act depends on the answers to two questions: what role you play, and how risky the use is.

1. Are you a provider or a deployer?

A **provider** develops an AI system, or has one developed, and places it on the market or puts it into service under its own name. A **deployer** uses an AI system under its own authority in the course of its work. A council using a vendor's chatbot, a hospital using a triage tool and an SME using AI features in its HR software are all deployers. Most Irish organisations are deployers, and most of their duties come from Articles 4, 26, 27 and 50.

Be careful with the line. If you put your own name on a system, make a substantial modification, or change its intended purpose so that it becomes high-risk, you can become a provider, with the much heavier provider obligations.

2. Which risk tier does each use fall into?

TIER	WHAT IT MEANS	EXAMPLES IN IRISH ORGANISATIONS
Prohibited Article 5	Banned outright since February 2025.	Social scoring that leads to unfair treatment; emotion recognition of staff or students; scraping faces from the internet to build recognition databases; manipulative techniques that cause significant harm.
High-risk Annex III	Allowed, with strict requirements from 2 December 2027.	CV screening and promotion decisions; assessing eligibility for social welfare, housing supports or other public assistance; exam marking and admissions; prioritising emergency calls; credit scoring.
Transparency Article 50	Allowed, if people are told. Applies since August 2026.	Public-facing chatbots; AI-generated images, audio or video; AI-generated text published to inform the public on matters of public interest, unless a human has reviewed it and takes editorial responsibility.
Minimal	No specific obligations beyond AI literacy.	Drafting and summarising with general-purpose assistants; spam filters; translation; scheduling.

Not every system used in an Annex III area is high-risk. A provider can document that a system only performs a narrow procedural task or prepares work for a human assessment (Article 6(3)), but systems that profile people are always high-risk. As a deployer, ask the provider for that assessment in writing.

Who regulates AI in Ireland

A **distributed model**: existing regulators supervise AI in their own sectors (for example the Data Protection Commission, the Central Bank and Coimisiún na Meán), coordinated by the **AI Office of Ireland**, the single point of contact with the Commission. Nine fundamental-rights bodies, including IHREC and the Ombudsman, can request the documentation providers and deployers must keep.

Eight key factors to get right

These are the factors that decide whether an organisation is ready, and the ones a regulator, an auditor or an evaluation panel is most likely to ask about.

1 Know your role for every system

Provider, deployer, importer or distributor: the role sets the duties, and it can change if you rebrand, modify or repurpose a system.

> Record the role against each system in the inventory, and have it checked whenever a system is customised.

2 Build and keep an AI inventory

List every AI system in use, including features inside existing software and tools staff use without approval. For each one note the purpose, the owner, the supplier, the data it uses and the people it affects.

> Make the inventory a living register owned by a named person, updated whenever software is bought or renewed.

3 Classify by risk, and write the reasoning down

Check each use against Article 5 and Annex III. Where a provider says an Annex III use is not high-risk, keep its Article 6(3) assessment on file.

> Stop any prohibited use now. Mark high-risk uses for the December 2027 work plan.

4 Treat AI literacy as an obligation, not a course

Article 4 has applied since February 2025. After the Omnibus, providers and deployers must *take measures to support* the AI literacy of staff and others using AI on their behalf, taking account of their knowledge, experience and the context of use. No certificate is required, and no single level is mandated. National authorities have been able to supervise the measures since August 2026.

> Train by role, tailor to the systems people actually use, and keep a record of what was done and for whom. A one-off video for everyone will not show proportionate measures.

The everyday reality

Most of the AI in a typical Irish organisation is minimal-risk, but it is often invisible: assistants built into email, office and HR software, browser extensions and free tools staff use on their own initiative. You cannot classify what you haven't found, so the inventory comes first.

5 Human oversight needs names, authority and time

For high-risk systems, deployers must assign oversight to people with the competence, training, authority and support to do it (Article 26(2)). Those people must understand the system's limits, watch for automation bias, and be able to disregard, override or stop it (Article 14).

> Name the overseer for each high-risk use, write down when they must intervene, and protect their time. Oversight that nobody has time for is not oversight.

6 Be open about AI

Since 2 August 2026, people must be told when they are interacting with an AI system unless it is obvious, and deployers must disclose deepfakes. AI-generated text published to inform the public on matters of public interest must be disclosed unless a person has reviewed it and takes editorial responsibility.

> Label chatbots, add an AI notice to public-facing services, and set a house rule for AI-assisted publications.

7 Connect the Act to GDPR, not beside it

Most AI systems process personal data, so the Data Protection Commission remains a key regulator. A data protection impact assessment and an AI risk assessment ask overlapping questions. The Omnibus also created a narrow legal basis for processing special category data to detect and correct bias, with strict safeguards.

> Run one joined-up assessment where you can, and involve your Data Protection Officer from the start.

8 Put it in the contract

A deployer can only meet its duties if the provider gives it what it needs: instructions for use, information on accuracy and limits, logs, support for oversight, and cooperation when something goes wrong. These are easiest to secure before signature, when the buyer has the most leverage.

> Add AI Act requirements to tenders and renewals now, including for systems that will only become high-risk in December 2027. See page 9 for the questions.

Penalties, for scale

Up to €35 million or 7% of worldwide turnover for prohibited practices; up to €15 million or 3% for most other breaches; up to €7.5 million or 1% for misleading information to authorities. SMEs and small mid-caps pay the lower of the two. Under the Irish legislation as published, fines on public-sector bodies are capped at €1 million.

What public bodies must add

Public bodies carry extra duties because their decisions affect access to services and rights. Most apply to high-risk uses from 2 December 2027, and the preparation takes longer than a year.

DUTY	WHAT IT REQUIRES	ARTICLE
Fundamental rights impact assessment	Before first using a high-risk system, assess how it will be used, how often, who could be affected, the risks of harm, the oversight measures and what happens if risks materialise. Notify the authority of the results.	27
Register the use	Public authorities using high-risk systems must register that use in the EU database, and must not use a system that the provider has not registered.	26(8), 49
Tell people	People subject to decisions made or assisted by a high-risk system must be informed. Workers and their representatives must be told before a high-risk system is used in the workplace.	26(7), 26(11)
Explain decisions	People affected by a decision based on a high-risk system, with legal or similarly significant effects, can ask for a clear explanation of the role the AI played.	86
Keep logs and monitor	Keep the system's logs for at least six months, monitor its operation, and report serious incidents.	26(5), 26(6)
Legacy systems	High-risk systems already in use by public authorities must comply by 2 August 2030.	111(2)

Irish guidance to build on

- The Government's **Guidelines for the Responsible Use of AI in the Public Service**, with a decision framework, a Responsible AI Canvas and AI lifecycle guidance.
- The National Cyber Security Centre's **Secure AI Adoption** guidance (June 2026), covering security across the AI lifecycle.
- The Commission's **AI literacy questions and answers**, which set out what proportionate literacy measures look like.

The practical test

If a citizen asked tomorrow whether AI was used in a decision about them, could the organisation say which system, who oversaw it and why the decision was made? If not, that is the work programme.

An AI literacy programme that stands up

The Commission's guidance is clear that one size does not fit all. Measures should reflect what AI the organisation uses, whether it is a provider or a deployer, the risks of those systems, and what each group of people already knows.

AUDIENCE	WHAT THEY NEED TO BE ABLE TO DO
Everyone who uses AI	Explain in plain terms what AI is and is not; recognise hallucinations, bias and data-protection risks in everyday tools; know what may and may not be entered into an AI tool; know when to say that AI was used, and whom to ask.
Managers and decision-makers	Judge whether a proposed AI use is sensible and lawful; understand the risk tiers and the prohibited practices; know what the organisation's register, policy and approvals require.
System owners and overseers	Understand a specific system's purpose, limits and failure modes; monitor it; resist automation bias; know when and how to override or stop it; keep the records the Act requires.
Procurement and contract managers	Identify AI in what is being bought; ask the right questions of suppliers; write Act requirements into tenders and contracts; check what is delivered against what was promised.
Leadership and boards	Set the organisation's appetite for AI; assign accountability; understand the regulatory landscape, the penalties and the reputational stakes; ask for evidence that controls work.

What good evidence looks like

- A literacy policy that links roles to the systems they use.
- Content refreshed when systems, guidance or the law change, as they did in 2026.
- Practical exercises on the organisation's own tools, not only theory.
- Completion records by role, including contractors acting on the organisation's behalf.
- Short checks of understanding, especially for overseers.
- Accessible formats, and Irish-language versions where the organisation's obligations require them.

> **Responsible AI is the content; the Act is the frame**
>
> The strongest programmes don't teach the Act article by article. They teach people to use AI well in their own jobs: checking outputs, keeping a human in charge of decisions, protecting personal data and being open about AI use. The legal duties follow.

Ten questions to ask every AI supplier

Most of the information a deployer needs sits with the provider. Ask for it in the tender, score the answers, and write the commitments into the contract.

1. Is your product, or any feature in it, an AI system under the Act? Which features, and which models do they use?
2. In which risk tier do you place each intended use, and can we see your reasoning, including any Article 6(3) assessment?
3. Are you the provider? If you rely on a general-purpose AI model, who provides it and what documentation do you hold?
4. What are the instructions for use, the known limits and the measured accuracy for our use case?
5. How does the system support human oversight: can our staff see why it produced an output, and override or stop it?
6. What logs does it keep, for how long, and how do we get them?
7. How do you meet the transparency duties, for example chatbot disclosure and marking of generated content?
8. For high-risk uses: is the system registered, and what will you give us for our fundamental rights impact assessment?
9. How will you tell us about changes, incidents and corrective actions, and how quickly?
10. What training do you provide for our users and overseers, and is it included in the price?

A 90-day plan

Days 1–30	Appoint an accountable owner. Build the AI inventory, including tools in use without approval. Stop any prohibited practice. Issue a short interim AI use policy.
Days 31–60	Classify each use and record your role. Launch role-based AI literacy, starting with everyone who uses AI and the owners of riskier systems. Add AI questions to every live procurement and renewal.
Days 61–90	Label chatbots and AI-generated content. Map the December 2027 work for each high-risk use: overseers, logs, impact assessments, registration and supplier evidence. Report progress to leadership.

Fifteen-point readiness checklist

Tick what you can evidence today, not what you intend to do.

- ☐ A named person is accountable for AI across the organisation.
- ☐ We have an inventory of every AI system in use, including features inside existing software.
- ☐ We know whether we are provider or deployer for each system.
- ☐ We have checked every use against the prohibited practices.
- ☐ We know which uses are, or may become, high-risk under Annex III.
- ☐ We have a written AI use policy that staff have seen.
- ☐ AI literacy measures are role-based and tailored to our systems.
- ☐ We keep records of who has been trained, including contractors.
- ☐ Each high-risk use has a named, trained overseer with authority to intervene.
- ☐ Chatbots and AI-generated content are clearly labelled.
- ☐ Our DPO is involved and AI assessments join up with DPIAs.
- ☐ Tenders and renewals include AI Act questions and contract terms.
- ☐ We know what our suppliers will provide for oversight, logs and incidents.
- ☐ Public bodies: we have a plan for impact assessments, registration and explaining decisions by December 2027.
- ☐ Leadership receives regular reports on AI use and risk.

About Dealcraft Advisory

Dealcraft Advisory is a Dublin consultancy that directs and writes technology bids and helps organisations buy technology well. This paper was prepared by **Daksh Kumar**, Founder. In his most recent role at IBM he led AI governance across consulting bids, making sure the use of AI in each proposal was governed, explained and accountable. He has more than 30 years' experience in technology deals and is the author of *The Bid Director: Win the Deals That Matter*.

To discuss AI governance, AI literacy or buying AI, email **sales@dealcraftadvisory.ie** or visit **dealcraftadvisory.ie**.

This paper is general information, current to November 2026, and is not legal advice. The AI Act, the Digital Omnibus and Irish implementing legislation should be read in full, and specific decisions should be taken with legal advice. © 2026 Dealcraft Advisory, Malahide, Co. Dublin. CRO Business Name No. 793378.

Sources

1. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Official Journal of the EU, 12 July 2024. Articles 3, 4, 5, 6, 14, 26, 27, 49, 50, 86, 99, 111, 113 and Annex III.
2. Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 (Digital Omnibus on AI), 8 July 2026; in force 27 July 2026. Consolidated text and recitals via artificialintelligenceact.eu/ai-act-explorer/digital-omnibus.
3. European Commission, Digital Omnibus proposal, 19 November 2025; political agreement 7 May 2026 (reported by Winston Taylor, "AI Act rules on high-risk AI delayed as AI Digital Omnibus agreed").
4. European Commission, "AI Literacy – Questions & Answers", digital-strategy.ec.europa.eu (updated after the Digital Omnibus).
5. Department of Enterprise, Tourism and Employment, "EU Artificial Intelligence (AI) Act", enterprise.gov.ie, including the list of fundamental rights authorities.
6. Department of Enterprise, Tourism and Employment, "Publication of the Regulation of Artificial Intelligence Bill 2026", 17 June 2026.
7. Matheson, "A deep dive into the Regulation of Artificial Intelligence Bill 2026" (AI Office, market surveillance authorities, sanctions including the €1 million cap for public bodies).
8. Dublin Chamber, "Ireland Opens its AI Office", 27 August 2026 (Regulation of Artificial Intelligence Act 2026 signed 21 July 2026; AI Office of Ireland; CEO Paul Byrne).
9. Houses of the Oireachtas Library & Research Service, Bill Digest: Regulation of Artificial Intelligence Bill 2026, 8 July 2026.
10. Department of Public Expenditure, Infrastructure, Public Service Reform and Digitalisation, "Guidelines for the Responsible Use of AI in the Public Service", gov.ie.
11. National Cyber Security Centre, "Secure AI Adoption", 30 June 2026, ncsc.gov.ie; reported by the Law Society Gazette, June 2026.
12. S.I. No. 366 of 2025, designating market surveillance and notifying authorities under the AI Act (as cited by Matheson).

Free resources from Dealcraft Advisory

More practical guides for organisations buying technology, including *What Should a New Business System Cost?* and *Starting Your ERP Journey*, are free at **dealcraftadvisory.ie/insights**.