JUNE 15, 2026

When the regulators arrive: what the EU AI Act and GDPR actually require of your AI

The EU AI Act's big deadlines just slipped to 2027 — but the obligations that catch an Irish SMB are already live. A plain-language map of what you owe now.

By **Dave Taylor**

If you run an Irish SMB using AI, two rulebooks already apply to you: the EU AI Act and GDPR. The headline AI Act deadlines just slipped to 2027 — but the parts that actually catch a small business are live now or land this August. GDPR has applied since 2018 and changes nothing about that. Here is what you owe, in plain language.

The deadline moved. Your exposure didn't.

In the first part of this series we mapped how AI works and where it breaks. This part picks up where the technology hands off to the regulator. And the regulator just did something that confused a lot of people. On 7 May 2026, the EU agreed a "Digital Omnibus" that postponed the AI Act's high-risk obligations: the use-case rules in Annex III slid from August 2026 to 2 December 2027, a 16-month delay, with product-embedded high-risk systems pushed to 2028. The coverage read as "AI Act delayed," and a lot of owners filed it under not-my-problem-yet.

That is the wrong read, and it's an expensive one. The high-risk regime — the heavy testing, documentation and conformity-assessment machinery — was never the part most SMBs had to worry about, because most SMBs don't build hiring algorithms or credit-scoring engines. The parts that do apply to you sit outside that delay, and two of them are already in force. So an owner who reads "delayed to 2027" and closes the tab has just talked themselves out of obligations that were live before they heard the headline. The deadline that moved is not the deadline that binds you, and the gap between the two is where the avoidable fines live.

Which risk tier are you actually in?

The AI Act sorts every system into one of four buckets, and knowing yours saves you most of the panic. **Prohibited** practices — social scoring, workplace emotion recognition, untargeted face-scraping — are banned outright and carry the top fines. **High-risk** covers AI used in hiring, credit, education, medical devices and critical infrastructure; this is the tier with the delayed, demanding obligations. **Limited-risk** is where most SMB tools actually live: chatbots, AI writing assistants, anything that talks to a customer or generates content. **Minimal-risk** is everything else — your spam filter, your AI photo editor — with no specific obligations at all.

The practical takeaway: unless you are using AI to screen job applicants, score creditworthiness or make decisions with legal weight, you are almost certainly in the limited-risk tier. That tier has exactly one headline duty — transparency — and it is manageable. The fines tell you the EU's priorities: prohibited-practice breaches run to €35 million or 7% of global turnover, while transparency breaches cap at €15 million or 3%. Know your tier before you budget for compliance.

The obligations that are already live

Two duties apply to you right now, regardless of risk tier. The first is **AI literacy**. Since 2 February 2025, every organisation that deploys AI must make sure the staff using it actually understand it — its limits, its risks, what it should and shouldn't be trusted with. There is no mandatory course or certificate; the EU explicitly wants a proportionate approach, lighter for a small firm than a bank. But "proportionate" is not "optional." If your team is pasting client data into ChatGPT with no guidance, you are already out of step with a duty that has been in force for over a year.

The second lands this August. From 2 August 2026, Article 50 requires you to **tell people when they're dealing with AI**. A customer-facing chatbot has to disclose it's a bot. AI-generated content — synthetic images, voice, video — has to be marked as artificially generated. This applies even to open-source tools, and it applies to you as the deployer, not just the vendor who built the model. If you've quietly put an AI receptionist or content generator in front of customers, this is the line item to action before the summer.

GDPR is the rule that bites first

Here is the part the AI Act headlines bury: for most Irish SMBs, GDPR is the regulation that actually has teeth today, and the AI Act delay changes nothing about it. The Irish Data Protection Commission has already published guidance flagging the GDPR risks in using AI — feeding it more personal data than you need, using it for automated decisions, and using it for purposes the person never agreed to. The DPC has open inquiries into how Google and X trained their models. AI is firmly on its desk.

Three GDPR duties attach the moment your AI touches personal data. You need a **lawful basis** for putting that data through a model — usually legitimate interests, but only after you've actually assessed it. You need a **DPIA**, a data protection impact assessment, for any high-risk processing — and an AI system making decisions about people generally qualifies. And under Article 22, a decision with legal or similarly significant effect can't be left to the machine alone; a person has to be able to review it. That last one is the same human-in-the-loop principle we reached on engineering grounds in Part 1's "excessive agency" — the law and the architecture arrive at the same place.

What it adds up to

Strip out the noise and the SMB obligation list is short. Know which risk tier each tool sits in. Make sure your people understand the AI they use. Tell customers when they're talking to a bot or reading generated content. And wherever AI touches personal data, have a

lawful basis, do the DPIA, and keep a human on any decision that matters. None of that requires the high-risk machinery that just got delayed to 2027 — it requires an honest inventory of what you've already switched on.

That inventory is the same one Part 1 ended on, which is not a coincidence: good security and good compliance ask the identical first question — what AI have we actually deployed, and what can it touch? Most owners can't answer that in one sitting, because the tools crept in one signup at a time. Part 3 turns it into a concrete exercise you can run in an afternoon. If you'd rather not map your AI exposure and obligations alone, book a 30-minute working session and we'll walk your stack — technical and regulatory — together.