

OCTOBER 19, 2026

Who enforces the EU AI Act in Ireland? 15 regulators, one coordinating office

Ireland didn't build one AI regulator. It handed the AI Act to 15 it already had. Here's who enforces what, and what that means for your board.

By Dave Taylor



Ireland enforces the EU AI Act through 15 regulators it already had, not one new one. The Central Bank, the Workplace Relations Commission (WRC), the Data Protection Commission (DPC), Coimisiún na Meán and others each police AI inside their own sector, while the new AI Office of Ireland coordinates them, keeps a national register and will run the regulatory sandbox. For a large firm, AI Act enforcement in Ireland can arrive from several directions at once. Ministers used last week's International AI Summit at the RDS to announce the sandbox and a National AI Literacy Taskforce. Here is who does what, and what a board should do about it.

How does Ireland's distributed enforcement model work?

Ireland chose not to build a single AI regulator. In September 2025 the Department of Enterprise, Tourism and Employment named [15 national competent authorities](#) – the bodies that already regulate banking, broadcasting, workplaces, energy, medicines, transport and personal data – and gave each one the AI Act brief for its own patch. The Regulation of Artificial Intelligence Act 2026, [signed into law on 21 July](#), put that model on a statutory footing.

The detail sits in [S.I. No. 405 of 2026](#), made on 31 July, which assigns the high-risk uses in Annex III of the Act to named bodies. The WRC takes employment: CV screening, candidate ranking, performance monitoring. The Central Bank takes credit scoring and life and health insurance pricing. The DPC takes biometrics, law enforcement, migration and the administration of justice. The Health Service Executive (HSE) covers eligibility for public benefits and emergency-call triage, while the Commission for Communications Regulation (ComReg), the National Transport Authority (NTA) and the Commission for Regulation of Utilities (CRU) split critical infrastructure. Coimisiún na Meán polices the transparency rules for media and hosting services, and product regulators such as the Health and Safety Authority (HSA) keep the AI built into machinery and devices. An engineer would recognise the design straight away: a distributed control system, where every node already knows its own plant and nobody sits in a central control room.

What does the AI Office of Ireland actually do?

The AI Office – Oifig IS na hÉireann – is the hub of that network, and its job description has changed more than once. It is a statutory body, chaired by Mary Doyle with Paul Byrne as chief executive, and was due to be operational by 2 August. When the Bill was published in June, it dropped the Office's planned role as a market surveillance authority; in [William Fry's reading](#), the Office would "simply coordinate the national regulatory system", maintain a national AI register and run regulatory sandboxes with priority access for small and medium-sized enterprises (SMEs) and start-ups. Six weeks later, S.I. 405 formally designated it a market surveillance authority as well, though coordination remains its main job. In practice, the Office is the front door for the European Commission and for business, it keeps the register and it will run the sandbox. It does not replace your sector regulator, and it will not be the body that investigates your credit model.

The money tells you where the weight sits. [Budget 2027](#) gave the Office an extra €2.3 million. Its total for the year is now €3.8 million. That is a coordination budget, not an enforcement one; the Central Bank and the DPC will run investigations in their own sectors with their own teams and their own supervisory habits. The sandbox – a supervised space

where a firm can test an AI system with the regulator before it goes live – was still described as forthcoming when ministers announced it at the [14 October summit](#), alongside the literacy taskforce and a €21 million AI Startup Accelerator. SMEs get priority, but large firms can apply. For a bank or insurer with a borderline credit or pricing model, a sandbox conversation is a far cheaper first meeting with a regulator than an inspection.

What did the Digital Omnibus change about the deadlines?

The EU's Digital Omnibus on AI, Regulation (EU) 2026/1744, came into force on 27 July and moved the heaviest obligations back. High-risk systems listed in Annex III – the hiring, credit, insurance and public-service uses above – now face their obligations from 2 December 2027, and AI embedded in regulated products from 2 August 2028, as [Cuatrecasas summarises](#). That is a 16-month extension on the original August 2026 date, and it leaves a large firm about 13 months from now to complete the groundwork: an inventory, risk classification, human-oversight design, technical documentation and logging. We covered the original delay and what it means for smaller firms in [what the EU AI Act and GDPR require of your AI](#). For a corporate reader the point is narrower. The extension buys time to build the compliance machinery properly. It does not remove any of the work, and a conformity assessment for a credit or hiring model needs evidence that has to be generated while the system runs.

Two sets of duties did not move. Article 50 transparency obligations – telling people they are dealing with a chatbot, and marking AI-generated content – have applied since 2 August 2026, with a grace period to 2 December 2026 for generative systems already on the market to add machine-readable marking. The prohibited practices that have applied since February 2025, from social scoring to emotion recognition in the workplace, are joined by new bans on nudification apps from 2 December 2026, which the [Omnibus added](#). Now picture one firm with a customer-service chatbot, a marketing team using image generators and an HR function piloting interview analytics. That is three regulators' worth of exposure in a single organisation: Coimisiún na Meán on transparency wherever media or hosting services are involved, the WRC on hiring, and the DPC on the personal data underneath all of it. None of the three will see what the other two see.

Which regulator will reach your firm first?

For a large Irish firm, the answer depends on what the AI is used for. The same underlying model can sit under three different regulators depending on whether it screens CVs, prices a policy or answers a customer. Here is the map as S.I. 405 sets it out for the main high-risk categories, which is the table every risk committee should have on file:

- **Screens CVs, ranks candidates, monitors or evaluates staff** → Workplace Relations Commission
- **Scores creditworthiness, prices life or health insurance** → Central Bank of Ireland
- **Biometric identification or categorisation** → Data Protection Commission
- **Assesses eligibility for public benefits, triages emergency calls** → Health Service Executive
- **Safety component in telecoms, transport or energy networks** → ComReg, NTA, CRU
- **AI built into machinery, toys or medical devices** → Product regulators such as the HSA and the Health Products Regulatory Authority (HPRA)

The penalties come from the AI Act itself. Breaching a prohibited practice carries fines of up to €35 million or 7% of worldwide turnover, most other breaches up to €15 million or 3%, and supplying misleading information to a regulator up to €7.5 million or 1%, under [Article 99](#). Large firms face whichever figure is higher; SMEs pay whichever is lower. The Irish Act adds a procedural step that boards should know about: under the general regime, William Fry notes that a fine only takes effect once the High Court confirms it. That gives a firm a route to challenge the amount. It does not offer a route around the investigation, the document requests or the reputational cost of a published finding. The bigger practical risk is overlap. A bank using AI in both lending and recruitment answers to the Central Bank and the WRC on the high-risk side and to the DPC on data, each with its own supervisory culture.

Why don't firms know which rules apply?

The gap between concern and readiness is measurable. In the [William Fry Technology Report 2026](#), based on an Ipsos B&A survey of 412 senior decision-makers, 81% of large businesses ranked legal and regulatory compliance as their number one consideration when adopting technology such as AI. Yet 55% of businesses said they were not aware of which EU laws apply to their technology adoption, and only 11% of larger businesses reported a positive return on AI investment so far. Put those three numbers side by side and you get a recognisable profile: compliance is the top concern, nobody can name the

rulebook, and the payback hasn't arrived. That is the firm whose AI pilots sit in legal review for months while nobody owns the list of AI systems already running in finance, HR and customer service.

A formal AI policy is the cheapest fix, and the evidence for it is unusually clear. Research by [Trinity College Dublin and Microsoft Ireland](#) found that only 44% of Irish organisations have a formal AI policy, and those that do are ten times more likely to report major productivity gains — 30% against 3%. A policy on its own won't satisfy a regulator. What it does is force the questions every one of the 15 authorities will eventually ask: what AI do you run, who owns each system, which risk category it falls into, who can override it, and where the logs are kept. A firm that has written those answers down can respond to a supervisory request from its own records. A firm that hasn't will be reconstructing them under deadline, usually with outside counsel billing by the hour.

Where this leaves boards and their advisers

Ireland's model puts the weight on the firm. With 15 regulators and a coordinating office running on €3.8 million, nobody in the system will hold a complete picture of your AI estate except you. That is the engineer's read of any distributed control system: it works when every node reports accurately, and it drifts quietly when the signal map is out of date. For a board, the signal map is the AI inventory — every system, its use case, its Annex III status and the regulator that owns it. Medical-device firms, where much of [our own operating background](#) sits, learned the same lesson under EU product regulation: the regulator audits what you wrote down at the time. For advisers, this is where the demand is — AI governance, risk classification, conformity assessment and assurance across financial services, HR and consumer-facing uses. The summit itself leaned heavily towards industry: [Tech Policy Press counted](#) 37 of 45 speaking slots going to companies or business groups.

The date that matters is 2 December 2027, and 13 months is less time than it sounds for a firm with dozens of AI systems and no inventory. Classifying a system is quick once the facts are on the table. Gathering them is the slow part: the vendor tools nobody logged, the model a team fine-tuned in a pilot, the scoring rule buried in a spreadsheet macro. Start with the map: we set out a practical method in [how to map your AI exposure](#), and the technical weak points a regulator will probe are covered in [how AI works and where it breaks](#). Part 2 of this series looks at what [Budget 2027 did, and did not do, for firms investing in AI](#). If you want a second pair of eyes on your AI inventory before a regulator asks to see it, [book a 30-minute working session](#).