Vivid Risk



BOARD GUIDE

AI and the EU AI Act: what the board needs to know, decide and ask

For boards of organisations in the EU, Ireland and the UK that use AI

Most organisations will meet the AI Act as users of AI, not as builders of it. This guide sets out what the Act asks of them and when, what AI changes about security, and a practical way to get ready. It uses the dates as amended by the Digital Omnibus on AI, which moved the high-risk deadlines.

THE SHORT VERSION

Five things every board member should know

1 **The Act looks at each AI system, not at the organisation**

One company can hold a prohibited use, a high-risk use and a harmless one at the same time. Nobody can answer "are we covered?" without a list of what is in use.

2 **Your duties depend on your role for each system**

If you use AI under your own authority, you are a deployer, and that is where most organisations sit. Put your own name on a high-risk system, modify it substantially, or change its purpose so that it becomes high-risk, and you take on the provider's duties (Art. 25(1)).

3 **The high-risk deadlines moved**

The Digital Omnibus on AI, Regulation (EU) 2026/1744, in force since 27 July 2026, moved the high-risk rules to 2 December 2027 for the Annex III uses and 2 August 2028 for AI in regulated products. Material written before it may give earlier dates.

4 **Some of it already applies**

The prohibitions and the AI literacy duty have applied since 2 February 2025, and the transparency duties since 2 August 2026. Two further prohibitions apply from 2 December 2026. Waiting for the high-risk date leaves these uncovered.

5 **AI is a security and data question before it is a regulatory one**

Staff putting client data into unapproved tools, or AI assistants acting with a person's access rights, are risks whatever tier the Act gives the system. Your existing cyber and data-protection duties reach them too.

> **Who answers for it?** NIS2 (Art. 20) and DORA (Art. 5) speak to the management body directly. The AI Act does not use that term: its duties and fines fall on the organisation. The question still comes to the board, because the board answers for how the organisation is run and the fines are large enough to matter.

Four decisions for the next board meeting

1. Name one executive who answers for AI use across the organisation.
2. Commission a register of the AI in use, with an owner for each system.
3. Approve a short AI use policy: which tools are approved and what data may go into them.
4. Agree what comes back to the board, and how often.

General information, not legal advice. Whether a rule applies to a given system is a judgement for your organisation and its advisers.

REGULATION (EU) 2024/1689

What the Act covers, and when

CATEGORY	WHAT FALLS IN IT	WHAT FOLLOWS
Prohibited Art. 5	Includes social scoring; inferring emotions at work or in education (except for medical or safety reasons); predicting crime from profiling alone; scraping faces to build recognition databases.	Stop the use.
High-risk Art. 6, Annexes I and III	Includes hiring and managing workers; credit scoring; life and health insurance pricing; education admissions and assessment; access to essential services; biometrics; safety components of critical infrastructure; AI that is a safety component of, or is itself, a product needing third-party conformity assessment under EU product-safety law.	The full set of duties for your role. A narrow exception exists (Art. 6(3)), never where the system profiles people.
Transparency duties Art. 50	Systems that interact with people or generate content; emotion recognition; deep fakes. These can sit on top of a high-risk system.	People are told. Providers mark synthetic output.
Everything else	Spam filters, grammar checks and most everyday tools.	Under the Act, AI literacy (Art. 4).

"Limited risk" and "minimal risk" are shorthand: the Act does not create categories with those names. Suppliers of general-purpose AI models have their own chapter (Arts. 53 and 55); simply using one does not make you its provider.

When it applies

2 February 2025	Prohibited practices and AI literacy.
2 August 2025	Rules for general-purpose AI models, and the penalties on operators.
2 August 2026	The Act generally applies, including the transparency duties. Providers' older generative systems have until 2 December 2026 to mark their output.
2 December 2026	Two further prohibitions, on non-consensual intimate imagery and child sexual abuse material.
2 December 2027	High-risk rules for the Annex III uses, including the deployer duties.
2 August 2028	High-risk rules for AI in products under EU product-safety law.

The fines the Act sets

Prohibited practices (Art. 99(3))	Up to €35,000,000 or 7% of worldwide annual turnover
Operator duties, including deployers (Art. 99(4))	Up to €15,000,000 or 3%
Incorrect, incomplete or misleading information to authorities (Art. 99(5))	Up to €7,500,000 or 1%

For an undertaking that is neither an SME nor a small mid-cap, the ceiling is the higher figure. For an SME or start-up it is the lower one on every line; for a small mid-cap, the lower one on the second and third lines (Art. 99(6) and (6a)).

Who it reaches, and who supervises

In the EU, Ireland included, the Act applies directly. **A UK organisation** is reached where it places AI systems on the EU market or where its systems' output is used in the EU, so the answer turns on where each system's output is used and whether you place systems on the EU market. Each Member State designates its authorities. **Ireland** has done so in S.I. No. 366 of 2025 as amended by S.I. No. 405 of 2026: for example the Workplace Relations Commission for employment, and the Central Bank of Ireland for regulated financial firms' credit and insurance systems.

Article references are to Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744, read in the Official Journal.

IF YOU USE AI

What the Act asks of a deployer

DUTY	IN PLAIN TERMS	ARTICLE
Every system, now		
AI literacy	Take measures to support the AI literacy of staff and others using AI for you. The Act sets no required level for any one person.	4
Stop prohibited uses	Check no use falls within Article 5.	5
Be open about AI content	Tell people exposed to emotion recognition or biometric categorisation. Disclose deep fakes, and AI-generated text on matters of public interest unless it had human review or editorial control.	50(3), (4)
High-risk systems: Annex III from 2 December 2027, regulated products from 2 August 2028		
Use it as instructed	Follow the provider's instructions for use.	26(1)
Human oversight	Give it to people with the competence, training and authority it needs.	26(2)
Input data	Where you control it, keep it relevant and sufficiently representative.	26(4)
Monitor and act	If it may present a risk, inform the provider or distributor and the authority, and suspend use.	26(5)
Keep the logs	Keep the logs under your control for at least six months, unless other law says otherwise.	26(6)
Tell workers first	Before using it at work, inform workers' representatives and the workers affected.	26(7)
Register your use	Public authorities register before using an Annex III system (not point 2), and do not use one that is unregistered.*	26(8), 49(3)
Connect it to your DPIA	Use the provider's information in any such assessment.	26(9)
Tell the people it decides about	Inform people that an Annex III system is used to make, or help make, decisions about them.	26(11)
Be ready to explain	A person affected by a significant decision based on an Annex III system can ask for a clear explanation of its role.*	86(1)
Fundamental rights impact assessment	Annex III systems, before first use: for public bodies and private bodies providing public services (except critical infrastructure), and for every deployer of credit-scoring or life and health insurance pricing systems. Then notify the authority of the results.	27(1), (3)

When something goes wrong. A deployer that identifies a serious incident informs the provider first, immediately, then the importer or distributor and the authority (Art. 26(5)). If it cannot reach the provider, Article 73 applies to it. That sets limits of 15 days, 10 days in the event of a death, and 2 days for a widespread infringement or a serious and irreversible disruption of critical infrastructure. An incident may also start clocks under NIS2, S.I. No. 360/2018, DORA or the GDPR, and those run separately.

If you build, or become the provider. A provider of a high-risk system carries the heavier load: a quality management system (Art. 17), technical documentation (Annex IV), conformity assessment (Art. 43), registration (Art. 49) and serious-incident reporting (Art. 73).

* Article 86 was not moved by the Omnibus, so on its face it applies from 2 August 2026, yet it concerns systems classified by rules that start later. Articles 49 and 73 raise the same question. Not confirmed as published when we read the sources on 14 October 2026: the fundamental rights assessment template, and final guidance on classification and on incident reporting.

Duties are summarised. Conditions and exceptions are in the articles cited.

THE SECURITY DIMENSION

AI changes what can go wrong

The Act's cybersecurity requirement is written for the provider: a high-risk system must be resilient against attempts to alter its use, outputs or performance, and Article 15 names data poisoning, model poisoning, adversarial inputs and confidentiality attacks. The organisation using the system still carries the risk of how it is connected, what it is given and who relies on it.

RISK	WHAT IT LOOKS LIKE	WHAT TO ASK MANAGEMENT
Unapproved tools ("shadow AI")	Staff use personal AI accounts or browser add-ons for real work, outside any review.	How do we find them, and is there an approved alternative that people actually want to use?
Sensitive data leaving	Client files, personal data or code pasted into a tool whose terms allow the supplier to keep or learn from it.	What data may go into which tool, and do the contract terms match?
Manipulated instructions	A document, email or web page carries hidden instructions that an AI assistant follows (prompt injection).	Which of our AI tools read outside content, and what can they do after reading it?
AI agents with real access	An agent acts with a person's permissions in email, files or finance systems, faster than anyone can check.	What can each agent do without a person approving it, and can we stop it?
AI switched on by a supplier	Software you already buy gains AI features through an update, and nobody reassesses it.	Who reviews supplier AI changes, and does our supplier review cover them?
Wrong answers relied on	A confident error reaches a customer, a decision or a filing, and no one checked it.	Where is a named person required to review the output before it is used?

Where they apply, your existing duties reach this

NIS2

EU, WHERE TRANSPOSED

AI tools in your network and information systems come within the risk-management measures (Art. 21) that the management body approves (Art. 20).

DORA

EU FINANCIAL ENTITIES

AI bought as an ICT service comes within ICT third-party risk, and the management body bears ultimate responsibility for ICT risk (Art. 5).

NIS Regulations 2018

UK · IRELAND

The UK is not under NIS2: its NIS Regulations and the NCSC CAF apply. Ireland has not yet transposed NIS2, so S.I. No. 360/2018 still applies, to operators of essential services and relevant digital service providers.

GDPR

EU · UK GDPR

Personal data put into AI needs a lawful basis and appropriate security. A data protection impact assessment may be required (Art. 35).

Good AI governance starts with what you already do for suppliers, access and incidents.

Technical defences for AI, such as data-loss prevention, access controls and adversarial testing, come from your security team and specialist providers.

GETTING READY

A six-step plan, and what the board should see

1 Find

List every AI system in use, including AI features inside software you already buy. Check procurement and expense records, single sign-on and app-consent logs and browser add-ons, and ask each team. Make it safe for people to say what they use.

Board sees: how many systems, and how the list was built.

2 Classify

For each system, record your role, the category, the reasoning, who decided and when. Where the answer is unclear, keep it open and say so rather than guessing. Revisit it when the use or the supplier changes.

Board sees: high-risk systems, and determinations still open.

3 Own

Give each system a business owner and, where high-risk, a person with the authority to oversee it. Legal, data protection, security, HR and procurement each hold part of this, so one executive has to hold the whole.

Board sees: a named owner for every system.

4 Control

Publish an approved-tools list and a short use policy. Add AI questions to supplier reviews: where data goes, whether it trains a model, what logs you can get, who to call. Limit what agents can touch. Define how a tool is approved, reviewed and retired. Train people for their role.

Board sees: training records, and unapproved tools found and dealt with.

5 Rehearse

Add an AI scenario to your incident exercise: a manipulated assistant, a data leak through a tool, a harmful decision. Practise who tells the provider, which reporting clocks start, and who can switch the system off.

Board sees: the exercise, what it exposed, and what was fixed.

6 Report

Bring AI into the regular risk report, beside cyber and data protection. Keep the records the Act asks for, so that the evidence exists before anyone requests it: oversight, logs, notices to workers and impact assessments.

Board sees: a short, regular report with the same measures each time.

Where to start. Steps 1 to 4 do not depend on any deadline, any guidance still to come or any purchase. Most of the value comes from a complete list and named owners. Start with the uses that touch hiring, credit, insurance, education and access to services.

Vivid Risk's view of a sensible order of work. The Act does not prescribe these steps.

Vivid Risk

Eight questions for your next meeting

1. Do we have a complete list of the AI we use, and how do we know?
2. Is any of it prohibited, and is any of it high-risk?
3. For each system, are we the deployer or the provider, and could that change?
4. Who owns each system, and who has the authority to stop it?
5. What data goes into these tools, and what do our contracts allow the supplier to do with it?
6. What have we done to support AI literacy, and can we show it?
7. If an AI system caused serious harm tomorrow, who would we tell, in what order, and by when?
8. What are we waiting for, and is it really a reason to wait?

How Vivid Risk can help, and where it stops

The AI Act workspace gives you

- A register of the AI systems you use, with owner, overseer and supplier.
- A guided path through the rules for each system, with the article behind each step.
- Your determination recorded, with reasoning, name and date.
- The duties for your role, cited and dated.
- Records for AI literacy, log retention and fundamental rights impact assessments, and a note of each notification you send.
- Article 73 clocks beside your other deadlines (with Compliance Workspace or Team).
- A register export, a dossier per system and an organisation summary.

What it does not do

- It does not classify your AI systems for you. The decision is yours.
- It does not find the AI in use. The register holds what you enter.
- It does not test, monitor or protect any AI system.
- It does not file, register or notify anything with anyone.
- It does not certify anything or carry out a conformity assessment.
- It lists a provider's duties and does not do the provider workload.
- It does not give legal advice.

Run the free scope check at vividrisk.ai/ai-act/scope-check, or keep your most important AI system in the workspace at no cost: the free snapshot covers one system, with no card.

info@vividrisk.ai
vividrisk.ai/ai-act

Vivid Risk LTD · Company registration number 818204 (Ireland)
Based in Dublin, Ireland

Written November 2026.
General information, not legal advice.