Prepared for

Leaders across every sector

Prepared by

itaa.ai

ITAA.AI WHITE PAPER · AI GOVERNANCE & REGULATION

The EU AI Act: What Every Organisation Needs to Know

A plain-English guide to the world's first comprehensive AI law: what it regulates, who it applies to, when the obligations bite, and where organisations will need help to be ready.

Return to the interactive online guide →

June 2026 • Updated for the Digital Omnibus reset • For general guidance, not legal advice

Contents

PART ONE – THE ESSENTIALS

Executive summary

THE HEADLINE

The EU AI Act is now law, it reaches far beyond Europe, and the recent delay is breathing space, not a reprieve.

The Act is the world's first comprehensive rulebook for artificial intelligence. It applies to organisations of every size and sector, including many based outside the EU, and it sorts AI by risk rather than by industry. The most serious obligations were eased in timing by the May 2026 Digital Omnibus, but the direction of travel is fixed.

The organisations that fare best will not be the ones that wait. They will be the ones that know which AI they already use, classify it honestly, and build the governance to manage it well.

Most leaders fall into one of two camps on the EU AI Act. The first assumes it is a problem for large technology companies and does not apply to them. The second knows it applies but assumes the recent delay means there is nothing to do yet. Both are mistaken, and both positions carry real cost.

This paper sets out, in plain language, what the Act does, who it binds, and what changes after the Digital Omnibus reset agreed in May 2026. It is written for general counsel and compliance leads, but also for the chief executive, the operations director, the head of HR, and the marketing lead, because the Act lands on all of them.

The Act takes a risk-based approach. A small number of AI uses are banned outright. A defined set of higher-stakes uses, many of them ordinary business activities such as recruitment screening or credit decisions, carry heavy obligations. A broader band of uses carry lighter transparency duties, such as telling people they are talking to a machine. And the large foundation models that sit underneath most modern AI tools have their own dedicated regime.

Three points deserve emphasis at the outset. First, scope is wide: if your AI system is used in the EU, or its output is used there, the Act can apply even if you are not. Second, most organisations are not building AI, they are buying and using it, which makes them deployers with their own real duties. Third, the penalties are GDPR-scale, reaching up to 35 million euros or 7 per cent of worldwide turnover for the most serious breaches.

TOP FINE	HIGH-RISK LIVE	IN FORCE NOW	REACH
7%	Dec 27	2	Global
Of global annual turnover, or 35m euros, for banned uses	New deadline for standalone high-risk systems	Banned uses and general-purpose AI rules apply now	Applies to non-EU organisations whose AI touches the EU

The honest message is that the delay helps, but it does not remove the work. The new deadlines give organisations time to do the job properly rather than in a panic. The job itself, knowing your AI, classifying it, governing it, and building the literacy to use it well, is exactly the work that makes AI deliver value in the first place. That is the through-line of this paper: compliance and good AI practice are the same discipline seen from two angles.

1. Why this matters to every organisation

It is tempting to read "EU AI Act" and conclude it is someone else's problem: a matter for Brussels, for big tech, or for the legal department. That instinct is wrong on every count, and the reasons are worth spelling out.

It reaches beyond the EU's borders

The Act applies to providers that place AI systems on the EU market, wherever they are based. It applies to organisations using AI systems within the EU. And, critically, it applies to providers and users outside the EU when the output produced by the system is used in the EU. A firm in London, New York, or Sydney that screens EU job applicants, serves EU customers through an AI chatbot, or sells software with AI features into Europe is squarely in scope. This is the same extraterritorial logic that made GDPR a global standard rather than a regional one.

It is about use, not industry

The Act does not regulate "the AI sector". It regulates what AI is used to do. The same large language model is unregulated when it drafts a marketing email and high-risk when it filters CVs for a shortlist. This means there is no industry that sits outside the Act. A charity, a professional association, a manufacturer, a law firm, and a hospital are all caught the moment they use AI for a regulated purpose.

You are probably already using regulated AI

Most organisations underestimate how much AI they already run, because much of it arrives embedded in software they already own. Recruitment platforms that rank candidates, tools that score creditworthiness, systems that monitor staff productivity, and customer-facing chatbots are all common, and several of these uses are either high-risk or carry transparency duties. The first surprise in almost every readiness review is the length of the list.

> **85%** of AI initiatives stall at the pilot stage, before delivering value. The same organisational weaknesses that stall AI projects also leave organisations exposed to the Act: no inventory, no ownership, no governance.

The cost of getting it wrong is material

Fines are set at GDPR scale and above. But the harder costs are often reputational and operational: an AI hiring tool found to discriminate, a chatbot that misleads customers, or a regulator's enquiry that freezes a product launch. For organisations whose credibility is their core asset, professional bodies, advisers, and public institutions among them, those costs dwarf the fine.

> **The reframe.** The Act is not only a compliance burden. It is a forcing function that pushes organisations to do what good AI adoption requires anyway: understand where AI is being used, decide who is accountable, and put guardrails around the decisions that matter. Treat it as governance, not paperwork, and the same effort that satisfies the regulator also makes the AI work.

2. What the EU AI Act actually is

The EU AI Act is a regulation, which means it applies directly across all EU member states without each country having to pass its own version. It entered into force on 1 August 2024 and applies in phases over the following years. It is the first law anywhere to regulate artificial intelligence comprehensively, and it is widely expected to shape rules well beyond Europe, just as GDPR did for data protection.

The core idea: regulate by risk

Rather than treating all AI the same, the Act sorts AI uses into tiers according to the risk they pose to people's safety, rights, and livelihoods. The greater the potential for harm, the heavier the obligations. A spam filter and a system that decides who gets a mortgage are both AI, but the law treats them very differently, and rightly so.

This produces a layered structure. A small number of uses are judged so harmful that they are banned. A defined set of higher-stakes uses are permitted but tightly controlled. A wider set carry only transparency duties. And the great majority of everyday AI uses carry no specific obligations at all beyond a general expectation of basic AI literacy among the people deploying them.

What counts as an "AI system"

The Act uses a deliberately broad definition, aligned with the OECD's: a machine-based system that, for explicit or implicit objectives, infers from the input it receives how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. In practice this captures modern machine learning and generative AI, while aiming to exclude simple, rule-based software. The breadth is intentional, and it is why an honest inventory matters: the boundary is wider than most teams assume.

Two regimes in one law

It helps to see the Act as containing two connected regimes. The first governs AI systems by their use, through the risk tiers described above. The second governs general-purpose AI models, the large foundation models such as those behind well-known chatbots, regardless of what they are eventually used for. The Digital Omnibus of May 2026 adjusted the timing and some of the detail of the first regime, while leaving the model rules largely as they were.

> **A moving target, handled carefully.** This paper reflects the position as at June 2026, including the Digital Omnibus political agreement of 7 May 2026. Some details remained subject to formal adoption, expected around July 2026. Where a figure or date could still shift, we say so. Always confirm the current position before making a compliance decision.

3. The four risk tiers, plus general-purpose AI

The risk-based structure is the heart of the Act. Understanding which tier a given AI use falls into is the single most important step in working out what you must do. The table below summarises the tiers; the sections that follow take the two demanding tiers in detail.

Tier	What it covers	What it means for you
Unacceptable risk (prohibited)	A short list of banned uses, including social scoring, manipulative or exploitative systems that cause harm, untargeted facial-image scraping, and emotion recognition at work or school.	Do not build, sell, or use these. The ban is already in force and carries the highest fines.
High risk	Permitted but heavily regulated uses in areas such as recruitment, credit, education, essential services, critical infrastructure, biometrics, and justice.	The bulk of the compliance work sits here: risk management, data governance, human oversight, documentation, and registration.
Limited risk (transparency)	Systems that interact with people or generate content: chatbots, deepfakes, and AI-generated media.	Tell people clearly when they are dealing with AI, and mark AI-generated content so it can be recognised.
Minimal risk	The large majority of AI uses: spam filters, recommendation engines, productivity tools, and the like.	No specific obligations beyond general good practice and basic AI literacy among staff.
General-purpose AI (separate regime)	The large foundation models that underpin many AI tools, with extra duties for the most capable models that pose systemic risk.	Mainly a duty on model providers, but it shapes the assurances buyers should demand from vendors.

How to read the tiers in practice

Two habits prevent most mistakes. The first is to classify the use, not the tool. Asking "is our chatbot high-risk?" is the wrong question; the right one is "what is this chatbot used to do, and does any of that fall into a regulated category?" A single tool can sit in more than one tier depending on how it is used. The second is to start from the banned list and work down, because confirming you are clear of the prohibitions is quick and removes the most serious exposure first.

> *The same AI model is minimal-risk when it drafts an email and high-risk when it screens a job application. The Act regulates the decision, not the software.*
>
> ITAA.ai

The remaining sections of Part Two work through the demanding parts of the Act in turn: the prohibitions, the high-risk regime, the transparency duties, the general-purpose AI rules, the revised timeline, the penalties, and the question of which role you occupy and therefore which duties you carry.

PART TWO – THE DETAIL

4. The red lines: prohibited practices

A small number of AI uses are judged to carry risks so unacceptable that they are banned across the EU. This prohibition has applied since 2 February 2025, making it the first part of the Act to bite. The bans sit in Article 5, and they attract the highest penalties in the whole regime.

What is banned

The prohibited uses include:

- **Harmful manipulation.** Systems that use subliminal, deceptive, or manipulative techniques to distort behaviour in ways that cause significant harm.
- **Exploiting vulnerability.** Systems that exploit vulnerabilities related to age, disability, or social or economic circumstances to materially distort behaviour and cause harm.
- **Social scoring.** Evaluating or classifying people over time based on behaviour or characteristics, leading to detrimental or disproportionate treatment.
- **Predictive policing of individuals.** Assessing the risk that a person will commit a crime based solely on profiling or personality traits.
- **Untargeted facial scraping.** Building or expanding facial-recognition databases by scraping images indiscriminately from the internet or CCTV.
- **Emotion recognition at work and school.** Inferring emotions in the workplace or in education, except for narrow medical or safety purposes.
- **Sensitive biometric categorisation.** Categorising people by biometric data to infer race, political opinions, trade union membership, religious beliefs, sex life, or sexual orientation.
- **Real-time remote biometric identification** in public spaces for law enforcement, except in tightly defined and authorised circumstances.

New prohibitions from the Digital Omnibus

The May 2026 reform added a further prohibition, taking effect on 2 December 2026: AI systems designed to generate or manipulate non-consensual intimate imagery, including so-called "nudifier" applications, and the generation of child sexual abuse material. This addition signals that the prohibited list is not closed; it will be extended as new harms emerge.

> **What most organisations should take from this.** The majority of organisations will not be doing any of these things deliberately. The risk is doing one of them unknowingly through a purchased tool, for example an employee-monitoring product that infers mood, or a marketing system that crosses from persuasion into manipulation of vulnerable groups. The action is to check, in writing, that nothing in your AI estate falls foul of the bans. It is a fast check that removes your single largest source of exposure.

35m euros or 7 per cent of worldwide annual turnover, whichever is higher, is the maximum fine for breaching the prohibitions. This is the most severe penalty tier in the Act.

5. High-risk systems: the heavy-lifting tier

This is where most of the Act's substantive obligations live, and where most organisations will find their real work. A use classified as high-risk is permitted, but only if a demanding set of requirements is met throughout the system's life. Crucially, many high-risk categories are ordinary business activities, not exotic technology.

What counts as high-risk

There are two routes into the high-risk tier. The first is AI used as a safety component of products already regulated under EU law, such as medical devices, machinery, vehicles, and lifts (these sit in what the Act calls Annex I). The second, and the one most organisations will meet, is a defined list of use cases set out in Annex III:

Area	Typical examples
Employment and workers	CV screening and candidate ranking, targeted job adverts, promotion and termination decisions, and monitoring or evaluating performance.
Essential services	Credit scoring and creditworthiness, risk assessment and pricing in life and health insurance, and eligibility for public benefits.
Education	Admissions and assessment, scoring of exams, and detecting prohibited behaviour during tests.
Biometrics	Remote biometric identification, biometric categorisation, and emotion recognition (where not already banned).
Critical infrastructure	Safety management of utilities such as water, gas, electricity, and digital infrastructure and traffic.
Law, justice, migration	Use by courts and in democratic processes, law enforcement, and migration, asylum, and border control.

The two categories most organisations encounter are **employment** and **essential services**. Any business that uses software to screen, rank, monitor, or evaluate staff, or to assess customers for credit or insurance, should assume it may be operating a high-risk system until it has confirmed otherwise.

The obligations on a high-risk system

Where a system is high-risk, the provider must build and maintain a substantial compliance apparatus, and the organisation using it carries duties too. In summary, a high-risk system requires:

- **A risk management system** that runs across the whole life of the product.
- **Data governance**, so that training, validation, and testing data are relevant, representative, and as free of error and bias as possible.

- **Technical documentation** and **automatic record-keeping** (logs) demonstrating compliance and enabling traceability.
- **Transparency** to the organisations deploying the system, with clear instructions for use.
- **Human oversight**, designed so that a person can understand, intervene in, and override the system.
- **Accuracy, robustness, and cybersecurity** appropriate to the system's purpose.
- **Conformity assessment** before the system goes on the market, registration in an EU database, and ongoing post-market monitoring.

The deployer's duties matter too

If you buy and use a high-risk system rather than build it, you are a deployer, and the Act gives you specific responsibilities. These include using the system in line with the provider's instructions, assigning competent human oversight, monitoring how it performs and reporting serious incidents, keeping the logs the system generates, and, in many cases, informing the people affected by its decisions. Public bodies and some private deployers must also carry out a fundamental rights impact assessment before use.

> **Where help is usually needed.** Few organisations have, today, the inventory, classification, and governance to know which of their systems are high-risk, let alone the documentation and oversight to satisfy the obligations. This is the largest single readiness gap we see, and it is rarely a technology problem. It is a question of organisational logic: who owns each system, who decides, and how oversight actually works in practice. We return to this in Part Three.

6. Transparency duties: chatbots, deepfakes and watermarks

Below the high-risk tier sits a band of uses that are not dangerous enough to be heavily regulated, but where people have a right to know that AI is involved. These transparency obligations, set out in Article 50, are light by comparison with the high-risk regime, but they touch a very large number of everyday systems, which is why they matter to almost everyone.

The four core duties

- **AI interaction.** When a person interacts with an AI system such as a chatbot, they must be told they are dealing with a machine, unless it is obvious.
- **Synthetic content.** AI-generated or manipulated audio, image, video, and text content must be marked in a machine-readable way so it can be detected as artificial.
- **Deepfakes.** Image, audio, or video content that has been generated or manipulated to resemble real people, places, or events must be clearly disclosed as such.
- **Emotion recognition and biometric categorisation.** Where these are used (and not otherwise banned), the people exposed to them must be informed.

What the Digital Omnibus changed

The reform adjusted the timetable rather than relaxing the duties. Specifically, the grace period for marking AI-generated content was reduced from six months to three, and marking requirements for systems already on the market apply from 2 December 2026. The other transparency duties, including chatbot and deepfake disclosure, still apply from 2 August 2026 without a transition. The practical effect is that any organisation deploying generative AI in customer-facing or public settings should plan its disclosure and labelling now, not later.

> **A common blind spot.** Marketing, communications, and customer-service teams are adopting generative AI fastest, often without legal sign-off. Auto-generated images, synthetic voice, AI chat, and AI-written content aimed at the public can all trigger transparency duties. A simple internal standard for disclosure and labelling prevents most problems and protects the organisation's credibility at the same time.

7. General-purpose AI and foundation models

The second regime within the Act governs general-purpose AI models: the large, flexible foundation models that can be adapted to a wide range of tasks and that underpin most of the AI tools organisations now use. These obligations have applied since 2 August 2025 and were largely untouched by the Digital Omnibus.

The baseline obligations on model providers

Providers of general-purpose AI models must maintain up-to-date technical documentation, provide information and documentation to the businesses that build on their models, put in place a policy to respect EU copyright law, and publish a sufficiently detailed summary of the content used to train the model. A voluntary Code of Practice exists to help providers show how they meet these duties.

Extra duties for models with systemic risk

The most capable models, those judged to pose systemic risk because of their scale and reach, carry additional obligations: model evaluation and adversarial testing, assessment and mitigation of systemic risks, incident reporting, and a higher standard of cybersecurity protection.

2 Aug 2027 Models already on the market before 2 August 2025 have until this date to come fully into line with the regime, giving providers a transition window.

Why this matters even if you only buy AI

Most organisations will never be a model provider. But this regime still shapes your risk, because you depend on these models through the tools you buy. The practical implication is in procurement: ask vendors which underlying models they use, whether those providers meet the general-purpose AI obligations, and what documentation and assurances they can pass on to you. A vendor that cannot answer is a risk indicator, not a neutral fact. Independent, vendor-neutral evaluation is one of the clearest places where outside help earns its keep.

8. The timeline: dates that have now moved

The Act applies in phases. By late 2025 it was clear that the high-risk timetable was running ahead of the technical standards and guidance needed to comply with it. In response, the European Commission proposed the Digital Omnibus in November 2025, and a political agreement was reached on 7 May 2026. The headline change is a significant deferral of the high-risk deadlines. The table below shows the position as at June 2026.

Date	What applies
1 Aug 2024	The Act enters into force, starting the phased clock.
2 Feb 2025	Prohibited practices are banned. The AI literacy duty begins to apply.
2 Aug 2025	General-purpose AI model obligations apply. Governance structures and the penalty provisions begin to apply (some enforcement powers phase in later).
2 Dec 2026 (new)	New prohibition on non-consensual intimate imagery and abuse material. Watermarking duties for in-market generative systems apply.
2 Aug 2027 (moved)	National regulatory sandboxes must be in place. Pre-2025 general-purpose models must be fully compliant.
2 Dec 2027 (moved)	Obligations for standalone high-risk systems (Annex III) apply. Originally 2 August 2026.
2 Aug 2028 (moved)	Obligations for AI embedded in regulated products (Annex I) apply. Originally 2 August 2027.

How to read the delay

The deferral is real relief, but it should be read as time to prepare properly, not permission to do nothing. Three things have not moved: the prohibitions, the general-purpose AI rules, and the AI literacy duty are all already live. And the high-risk work, building an inventory, classifying systems, and standing up governance, takes many months when done well. An organisation that starts in 2027 will be doing in a panic what it could have done calmly across 2026.

A planning note. The Digital Omnibus also eased the rules for smaller organisations. The simplified compliance framework, with lighter documentation, reduced fines, and sandbox access, is being extended to a new "small mid-cap" category: organisations with fewer than 750 employees and turnover not exceeding 150 million euros. Many mid-sized organisations that assumed they were too large for relief now qualify. This is worth confirming for your own organisation.

9. Penalties and enforcement

The Act backs its obligations with penalties on a par with, and in places exceeding, GDPR. Fines are tiered by the seriousness of the breach, and for each tier the cap is the higher of a fixed sum or a percentage of worldwide annual turnover.

Breach	Maximum fine	Examples
Prohibited practices	**35m euros or 7% of turnover**	Using or supplying a banned AI system.
Most other obligations	**15m euros or 3% of turnover**	Breaching high-risk duties, transparency duties, or the obligations of providers, deployers, importers, and distributors.
Incorrect information	**7.5m euros or 1% of turnover**	Giving incorrect, incomplete, or misleading information to authorities or notified bodies.

For small and medium-sized enterprises and start-ups, the fine is the lower of the fixed sum and the percentage, rather than the higher, which softens the impact on smaller organisations. Penalties are required to be effective, proportionate, and dissuasive, and to take account of the organisation's size and circumstances.

Who enforces the Act

Enforcement is shared. A central European AI Office oversees the general-purpose AI regime and coordinates consistent application across the EU. Each member state designates national authorities to supervise and enforce the rest of the Act within its territory. The Digital Omnibus clarified the division of responsibility, giving the AI Office clearer competence over certain general-purpose systems and over AI built into the largest online platforms, while national authorities retain oversight of areas such as law enforcement and financial services.

The real exposure is rarely the headline fine. For most organisations, the more likely consequences are an investigation that stalls a product, a contractual dispute when an AI tool fails to meet the Act, or reputational damage when an AI decision is challenged. Good governance is cheaper than any of these, and it is the same governance that makes AI deliver value.

10. Who is responsible: your role under the Act

The Act assigns duties according to the role you play in relation to an AI system. The same organisation can hold different roles for different systems, so the first step is to map each significant system to a role. The four main roles are below.

Role	Who you are	Core responsibility
Provider	You develop an AI system, or have one developed, and place it on the market or put it into service under your own name.	The heaviest duties, especially for high-risk systems: conformity, documentation, and registration.
Deployer	You use an AI system in the course of your activities. This is the role most organisations occupy.	Use as instructed, ensure human oversight, monitor, keep logs, and inform affected people.
Importer	You place on the EU market an AI system from a provider based outside the EU.	Verify the provider has met its obligations before the system is sold.
Distributor	You make an AI system available on the market without being the provider or importer.	Check that required markings and documentation are in place.

The catch most organisations miss

Most organisations are deployers, and assume that the bulk of the responsibility therefore sits with the vendor. It does not. Deployers carry their own real duties, and there is an important trap: if you significantly modify a high-risk system, or put your own name on it, or use it for a purpose that makes it high-risk, you can become a provider in the eyes of the Act, with all the heavier obligations that follow. Fine-tuning a model, rebadging a tool, or building a bespoke workflow on top of a bought system can all trigger this. Knowing your role is therefore not a formality; it determines the size of the job.

> *Buying the AI does not outsource the accountability. Under the Act, the organisation that uses a system to make decisions about people owns duties of its own.*
>
> ITAA.ai

PART THREE – WHAT IT MEANS FOR YOU

11. AI literacy: the duty that is already live

ALREADY IN FORCE · APPLIES TO ALMOST EVERYONE

Of every obligation in the Act, AI literacy is the one that already applies, to nearly every organisation, and the most sensible place to begin.

Article 4 requires organisations to make sure the people who use AI on their behalf actually understand it. It has applied since 2 February 2025, alongside the prohibitions, and it is not limited to high-risk systems. If your staff use AI, this duty is yours today.

The Act asks for a "sufficient level" of AI literacy among staff and anyone operating AI on the organisation's behalf. "Sufficient" is deliberately contextual: it scales with a person's role, the systems they use, and the people affected by them. A board member, a recruiter using a screening tool, and a marketer using generative AI each need a different level and kind of understanding. The Digital Omnibus softened the wording from a duty to ensure literacy to a duty to support its development, but the expectation, and the exposure if it is ignored, remain.

Crucially, literacy is the foundation the rest of the Act stands on. The human oversight that high-risk systems require is only real if the people doing the overseeing understand what they are looking at: you cannot meaningfully challenge or override a system you do not understand. Literacy is also the practical defence against the risks that land first, before any high-risk deadline: misuse, over-reliance, leaking data into public tools, and the "shadow AI" that staff adopt without sign-off.

What "good" looks like, by role

Effective literacy is not a single e-learning module sent to everyone. It is tailored to roles, tied to the organisation's actual AI uses and policies, and refreshed as the tools change. A workable shape:

- **Board and executive:** enough to ask the right governance questions, set risk appetite, and own accountability for AI decisions.
- **Managers:** enough to oversee AI-assisted decisions, recognise when to intervene, and apply policy consistently.
- **Frontline users:** enough to use tools safely, spot errors and bias, handle data responsibly, and avoid shadow AI.
- **Specialists and buyers:** enough to evaluate vendors, models, and data, and to document the decisions behind them.

Why start here. Literacy is the lowest-cost, highest-leverage move an organisation can make on the Act. It satisfies a duty that is already live, it reduces the risks that surface first, and it builds the shared understanding that every later step, inventory, classification, and governance, depends on. It is where most organisations should begin, and where ITAA.ai most often starts with clients: our AI Foundations and AI Training & Enablement programmes build role-tailored literacy mapped to your actual AI estate and obligations.

12. The implications, function by function

The Act is often discussed as a legal matter, but its obligations land in operational functions across the organisation. Here is where the duties most commonly surface, and what each function should be asking.

Human resources and people

HR is the function most exposed, because recruitment, performance management, and workforce monitoring are explicitly high-risk. Any tool that screens, ranks, scores, or monitors people needs to be identified, classified, and governed, with human oversight that is real rather than nominal. Candidates and staff may need to be informed. This is the first place most organisations should look.

Finance, credit, and insurance

Creditworthiness assessment and insurance risk and pricing are high-risk uses. Finance teams and insurers using AI for these decisions face the full set of obligations, alongside existing financial regulation. Documentation, bias testing, and explainability are central here.

Marketing and communications

The transparency duties bite hardest here. Chatbots, AI-generated imagery and copy, synthetic voice, and deepfake-style content all carry disclosure or labelling duties. The function also carries the manipulation risk: persuasion is legitimate, but systems that exploit vulnerability are not.

Operations and IT

Operations and IT typically own the AI estate in practice, even when individual tools are bought by other teams. They are usually best placed to build and maintain the inventory, manage vendors, and operate the logging, monitoring, and security the Act expects. Shadow AI, tools adopted by staff without sign-off, is their hardest problem.

Legal, risk, and compliance

This function owns classification, the assessment of obligations, incident reporting, and the fundamental rights impact assessment where required. It cannot do the job alone, because it rarely has full visibility of what is in use, which is why cross-functional ownership matters.

The board and executive

Accountability ultimately sits at the top. The board should expect a clear answer to three questions: what AI are we using, where are we exposed, and who owns the response. An inability to answer is itself the finding.

> **The pattern across all of these.** Notice that the Act does not respect organisational silos. The same system may concern HR, IT, legal, and the board at once. That is exactly why AI governance fails when it is treated as a single department's task. It is an organisational design question: how decisions are made, who is accountable, and how oversight works across functions.

13. The readiness gap: where help is needed

In our work with organisations preparing for the Act, the same gaps recur. They are rarely about technology, and almost always about organisational logic: the structures, decisions, and ways of working that determine whether AI can be governed at all. These are the places where most organisations will need support.

You cannot govern what you cannot see

The first gap is visibility. Most organisations have no complete inventory of the AI they use, partly because so much of it is embedded in other software and partly because staff adopt tools informally. Without an inventory, classification is impossible and every later step is guesswork.

Classification is harder than it looks

Deciding which tier a system falls into, and which role you hold, requires judgement about how a system is actually used, not just what it is. This is where organisations most often get it wrong, either over-classifying and wasting effort, or under-classifying and carrying hidden exposure.

Governance that exists on paper only

Many organisations have an AI policy. Far fewer have working governance: named owners, a route for sign-off, real human oversight, incident reporting that functions, and a board that sees the right information. The Act demands the working version, not the document.

Literacy is now a legal duty, not a nicety

The Act requires organisations to ensure a sufficient level of AI literacy among the people who deploy AI on their behalf. The Digital Omnibus softened the wording to a duty to support the development of literacy, but the expectation remains. Staff who do not understand the tools they use cannot provide the oversight the Act requires.

NO INVENTORY	POLICY VS PRACTICE	LITERACY	OWNERSHIP
Most	**Gap**	**Duty**	**Unclear**
Organisations cannot list the AI they use	A policy exists; working governance does not	Now an obligation, not an optional extra	No single accountable owner for AI risk

Each of these gaps is bridgeable, and none requires waiting for the 2027 deadlines. They are also, not coincidentally, the same foundations that allow AI to move from scattered experimentation to measurable value. Closing the readiness gap and getting value from AI are the same programme of work.

14. A practical roadmap to readiness

The path to readiness is not complicated, but it does have an order. The sequence below can be started now and completed well within the revised timeline. Each step builds on the one before.

1. **Build an AI inventory.** List every AI system in use across the organisation, including tools embedded in other software and those adopted informally by staff. This is the foundation for everything else.
2. **Check against the prohibitions.** Confirm in writing that nothing in the inventory falls into the banned list. This removes your highest-exposure risk quickly.
3. **Classify each system by tier and role.** For each system, decide its risk tier and your role (provider, deployer, importer, or distributor). This determines the size of the task for each one.
4. **Assess your high-risk and transparency obligations.** For high-risk systems, map the gap against the requirements. For customer-facing generative AI, plan disclosure and labelling now.
5. **Stand up working governance.** Assign named owners, a sign-off route, real human oversight, and incident reporting. Give the board the visibility it needs to be accountable.
6. **Build AI literacy.** Equip the people who use AI to understand it well enough to oversee it. Tailor the training to roles rather than running one generic course.
7. **Tighten procurement.** Make vendor assurances about the Act a standard part of buying any AI-enabled tool, and document what you are told.
8. **Review and repeat.** The inventory and classification are living documents. Set a cadence to refresh them as tools, uses, and the rules themselves change.

> **The order matters.** Most failed compliance efforts start in the wrong place, usually by writing a policy before anyone knows what is actually in use. Inventory first, then classify, then govern. The policy is the last step, not the first, because it should describe how the organisation actually works, not how someone hopes it might.

15. How ITAA.ai can help

ITAA.ai helps organisations move from scattered, informal AI use to a governed, intentional approach that is both compliant and genuinely valuable. We are independent and vendor-neutral, governance-led, and human-centred, and our work starts from organisational logic rather than from technology. The Act is, for us, a particularly clear case of the same problem we exist to solve: making intelligence operational, responsibly.

Our services map directly onto the readiness gaps and the roadmap above.

Service	How it supports EU AI Act readiness
AI Assessments	An independent review of your AI estate, governance, and capability, including the inventory and risk classification that readiness depends on.
AI Strategy & Development	A purpose-led strategy and roadmap that build compliance into how AI delivers value, rather than treating it as a separate burden.
AI Implementation Planning & Integration	Translating classification and obligations into working governance: owners, oversight, sign-off, and incident reporting that function in practice.
AI Technology Evaluation & Vendor Selection	Independent, vendor-neutral assessment of tools and the assurances their providers can give you, strengthening procurement.
AI Foundations & AI Training & Enablement	Role-tailored AI literacy programmes for board, managers, and frontline staff that satisfy the Article 4 duty (already in force) and make human oversight real. For most organisations, this is the natural first engagement.
AI Advisory & Assurance	Ongoing oversight and continuous improvement as your AI use, and the rules, keep changing.

Our approach draws on three frameworks that keep the work grounded in how organisations actually function: the AI Strategy Pyramid, which connects purpose to strategy to execution; the Five-Layer Integration Model, which traces AI from back-end systems through to the public; and the Human-AI-Data Framework, which keeps people, technology, and information in balance. Each helps ensure that readiness for the Act is built on the organisation's real structures, not bolted on beside them.

THE BOTTOM LINE

Readiness for the EU AI Act is not a separate project. It is good AI governance, which is the same thing that makes AI deliver value.

The organisations that treat the Act as a forcing function, and do the foundational work now, will be both compliant and ahead. We would welcome a conversation about where your organisation stands and what a sensible first step looks like.

To discuss your organisation's readiness, contact ITAA.ai at **aking@itaa.ai**.

Want to learn more? Explore our full range of services and our approach to making AI work for organisations at **www.itaa.ai**.

Return to the interactive online guide →

Sources and further reading

This paper reflects publicly available information on the EU AI Act and the Digital Omnibus as at June 2026. The Digital Omnibus reached political agreement on 7 May 2026, with formal adoption expected around July 2026; some details remained subject to change at the time of writing.

- European Commission, *AI Act: Shaping Europe's digital future*. digital-strategy.ec.europa.eu
- EU AI Act Service Desk, *Timeline for the implementation of the EU AI Act*; *Article 99: Penalties*. ai-act-service-desk.ec.europa.eu
- Council of the EU, *Artificial intelligence: Council and Parliament agree to simplify and streamline rules*, 7 May 2026. consilium.europa.eu
- *EU AI Act Omnibus Agreement: postponed high-risk deadlines and other key changes*, Gibson Dunn.
- *EU AI Act update: timeline relief, targeted simplification, and new prohibitions*, Covington (Inside Privacy / Global Policy Watch).
- *Rules on high-risk AI to be delayed under EU omnibus deal*, Pinsent Masons (Out-Law).
- *EU agrees Digital Omnibus deal to simplify AI rules*, White & Case.
- Article 5 (prohibited practices), Article 50 (transparency), Articles 51 to 55 (general-purpose AI), and Annex III (high-risk use cases), EU AI Act. artificialintelligenceact.eu

Important. This white paper is general guidance, not legal advice. The EU AI Act is detailed and still settling, and its application depends on the specific facts of each organisation and system. Confirm the current rules and take professional advice before making compliance decisions.