AIRAS Cloud

THE IRISH GUIDE TO AI GOVERNANCE

What the EU AI Act means now -
and what Irish organisations should do next

Written by Richie Higgins
Founder & CEO, AIRAS Cloud

Media & Business Edition

3 August 2026

A NOTE FROM THE FOUNDER

Ireland does not need more AI panic. It needs an operating plan.

I wrote this guide because the conversation around the EU AI Act has become louder, but not always clearer. Irish businesses are hearing about enforcement, transparency duties and enormous fines. Many are still unsure where to begin.

My view is straightforward: do not start with a 100-page policy. Start with the AI that is actually being used. Find it. Name the owner. Determine the organisation's role. Screen what is prohibited. Identify the duties that apply now. Keep the evidence. Then build a repeatable governance process around it.

That is what this guide is designed to help you do. It is written for boards, business owners, technology teams, compliance leaders, public bodies and Irish SMEs that need practical direction rather than another legal summary.

The principle behind this guide

AI can support discovery, assessment and workflow. The final accountable decision must remain with an authorised human. AI does not approve AI.

AIRAS Cloud is featured throughout because it is the platform I built to solve this operating problem. The guidance stands on its own. The product is there for organisations that want to put the guidance into practice efficiently, consistently and at a cost proportionate to their scale and risk.

Richie Higgins

Founder & CEO, AIRAS Cloud

AFRH Consulting Limited | CRO 798243 | Ireland

Published 3 August 2026. Regulatory position reviewed against official EU and Irish sources available on that date.

What changed - and what Irish organisations should do now

The short position

On 2 August 2026, Article 50 transparency obligations began applying and enforcement started for rules already in scope, including general-purpose AI, prohibited practices, transparency and AI literacy. The high-risk timetable is staged and now extends to 2 December 2027 and 2 August 2028.

1 Build an AI inventory Include embedded vendor AI, staff-adopted tools, APIs and agents.	**2 Assign named owners** Responsibility needs a person, not a department.
3 Determine your role Provider, deployer, importer or distributor - per system.	**4 Screen prohibited practices** Do this before ordinary risk scoring.
5 Assess Article 50 Identify interaction, labelling and marking duties.	**6 Build role-based AI literacy** Different roles need different competence.
7 Use one repeatable method Keep rules, evidence, assumptions and versions visible.	**8 Keep approval human** Separate assessor and approver where required.
9 Monitor material change Models, vendors, permissions and uses change.	**10 Preserve the evidence** Show what was known, decided, controlled and reviewed.

Source note: European Commission and EU AI Act Service Desk [1]-[3].

2 August 2026 is a major milestone - not the only deadline

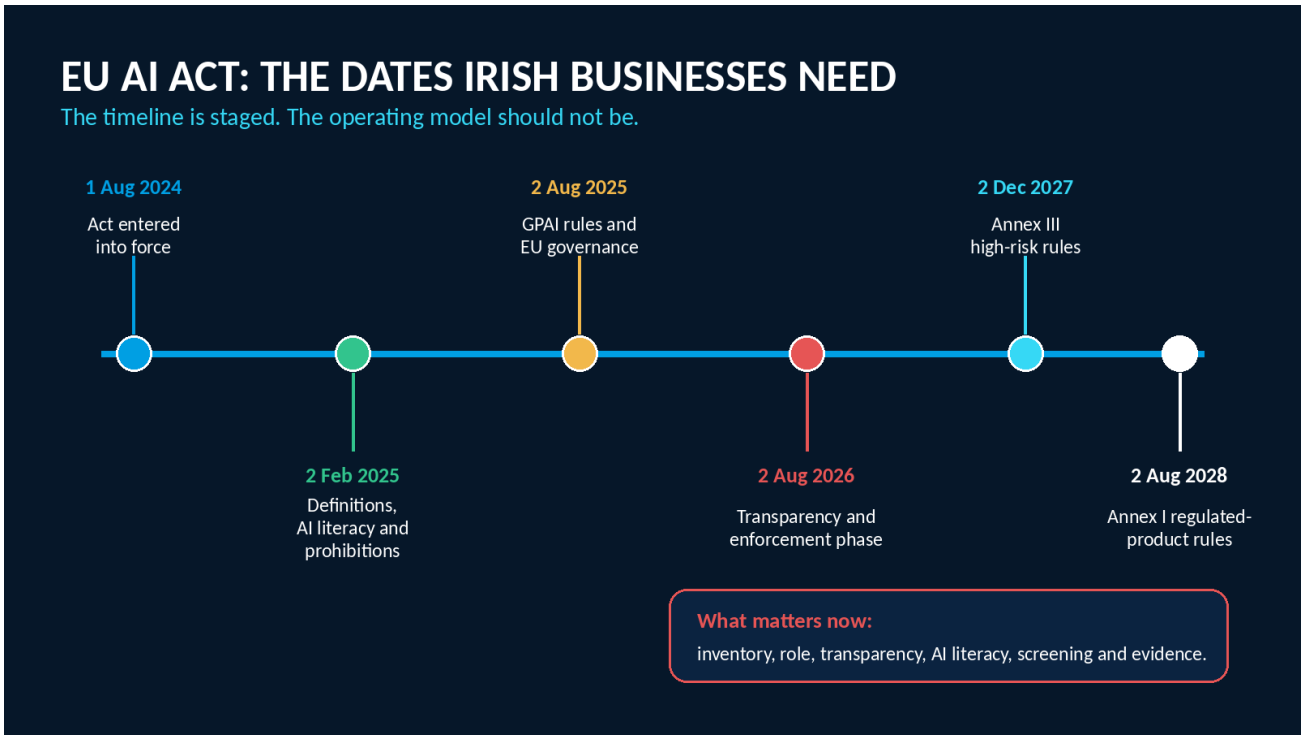

Official dates simplified for business planning. See References [1]-[3].

The EU AI Act applies progressively. Prohibitions and AI literacy began applying in February 2025. General-purpose AI model rules began in August 2025. The 2 August 2026 milestone brings Article 50 transparency obligations into application and starts enforcement for currently applicable rules.

The 2026 Digital Omnibus amended the timetable for high-risk systems. Annex III high-risk rules are scheduled for 2 December 2027, while high-risk AI embedded in products covered by Annex I is scheduled for 2 August 2028.

What I would not do

I would not tell every team that every future high-risk duty applies today. I would build the inventory, ownership, role analysis, transparency controls and evidence now so the organisation is ready for each staged deadline.

The EU rules are now backed by a domestic enforcement structure

The EU AI Act applies directly in Ireland. Ireland has also enacted the Regulation of Artificial Intelligence Act 2026, signed on 21 July 2026. The Irish Act provides the national enforcement architecture, including Oifig IS na hÉireann - the AI Office of Ireland - and a distributed model involving sectoral authorities. Commencement can be staged by ministerial order.

For an Irish organisation, that means AI governance will not sit with one regulator in every case. The relevant authority may depend on the sector, product, service and rights affected. Financial services, health, workplace, consumer, communications, data protection and product-safety issues may involve different supervisory expertise.

Practical implication

Do not create an “AI compliance” silo that ignores the systems already responsible for privacy, cybersecurity, quality, consumer protection, employment, procurement and product safety. AI governance must connect them.

The Irish Act does not replace the obligations in the EU Regulation. It gives Irish authorities the mechanisms to supervise, investigate and enforce them. That is why evidence quality now matters as much as policy wording.

Source note: Regulation of Artificial Intelligence Act 2026 (Act No. 31 of 2026) and Houses of the Oireachtas [4]-[5].

Know the role your organisation holds for each AI system

Role determination must be made per AI system and use case.

A company can be a deployer of a purchased chatbot, a provider of an internal model placed into a customer service, and an importer of a system supplied from outside the EU. The label changes the duties.

- Record the role per system, not once for the whole company.
- Preserve the evidence behind the role decision.
- Revisit the role when the system is materially modified, rebranded or placed on the market under your name.
- Do not assume a supplier contract transfers every responsibility away from your organisation.

Source note: Irish Department of Enterprise guidance and the EU AI Act role definitions [6].

Before you classify risk, confirm that you are dealing with an AI system

Not every piece of software is an AI system. Equally, a system does not need to learn continuously after deployment to fall within the Act. The definition focuses on a machine-based system that operates with varying levels of autonomy and may exhibit adaptiveness, inferring from inputs how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments.

A useful test

Ask whether the system goes beyond executing a fixed human-written rule and instead infers how to produce an output that influences a person, process or environment.

Systems that commonly deserve review include:

- Generative assistants and copilots
- Decision-support and scoring models
- Recommendation and ranking systems
- Computer vision and biometric systems
- Fraud, anomaly and predictive models
- Embedded AI features inside SaaS products
- Autonomous or semi-autonomous agents
- AI-enabled recruitment, credit, health or public-service tools

The safest approach is not to over-classify everything. It is to use a qualification process that records the evidence, what remains uncertain and who confirmed the conclusion.

Source note: EU AI Act definition guidance and Irish Department of Enterprise summary [6]-[7].

STEP TWO

Build one reliable AI inventory

Most organisations can name the AI projects approved by a steering committee. The real exposure is often wider: embedded vendor features, free tools, browser extensions, local software, APIs, pilots, shadow AI and agents that receive new permissions after approval.

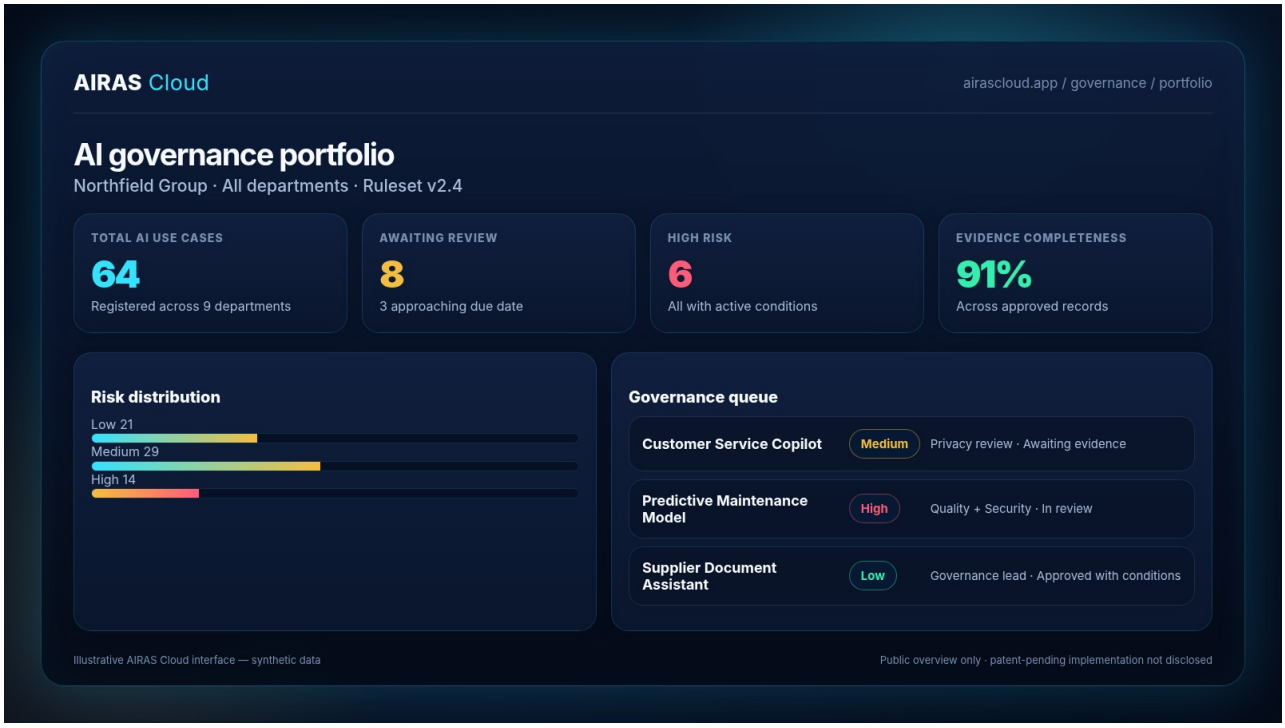

Illustrative AIRAS Cloud interface - synthetic data. A single portfolio view for ownership, risk, review and evidence status.

Minimum inventory fields

- ✓ System or capability name
- ✓ Business and technical owners
- ✓ Organisation role
- ✓ Data categories and sources
- ✓ Risk and transparency status
- ✓ Decision and conditions
- ✓ Business purpose and excluded uses
- ✓ Vendor and model provider
- ✓ Users and affected people
- ✓ Systems, APIs and write-back access
- ✓ Evidence completeness
- ✓ Next review date and change triggers

A register is not proof of completeness

The inventory should show how systems were discovered, what sources were checked and where coverage remains uncertain. A manual list alone cannot prove that embedded or shadow AI has been found.

Screen prohibited practices before ordinary risk scoring

Prohibited practices are not simply “very high risk”. They are a separate legal gate. A positive indicator should stop or escalate the use case before the organisation spends time debating controls or calculating an average score.

The screen should examine, where relevant:

- Manipulative or deceptive techniques causing significant harm
- Exploitation of vulnerabilities linked to age, disability or social/economic situation
- Certain social-scoring practices
- Certain predictive criminal-risk profiling
- Untargeted scraping of facial images
- Certain emotion-recognition uses in workplaces or education
- Certain biometric categorisation practices
- Real-time remote biometric identification in publicly accessible spaces, subject to narrow exceptions
- The additional prohibited content-generation practices introduced by the 2026 amendment from 2 December 2026

My rule for governance teams

Do not allow a prohibited-practice flag to be averaged down by ten lower-risk answers. The decision path must make hard stops visible and non-bypassable.

Source note: Article 5, Digital Omnibus amendments and implementation timeline [2]-[3], [8].

Transparency is now an operating control

From 2 August 2026, Article 50 requires certain providers and deployers to inform people when they are interacting with AI and to mark or disclose certain AI-generated or manipulated content. The exact duty depends on the system and context.

Direct interaction Inform people when they are interacting with an AI system, unless it is obvious in context.	**Synthetic content** Providers of relevant systems must support machine-readable marking and detectability, subject to the legal conditions and exceptions.
Deepfakes and public-interest content Deployers may need clear disclosure for deepfakes and certain public-interest text without human editorial control.	**Emotion / biometric categorisation** People exposed to defined emotion-recognition or biometric-categorisation systems must be informed, subject to exceptions.

Questions to put into the assessment

- Is a person interacting directly with the AI?
- Could the interaction reasonably appear human?
- Does the system generate or manipulate image, audio, video or text?
- Is the output a deepfake within the legal definition?
- Is public-interest text published without human editorial responsibility?
- Who owns the disclosure, marking, testing and evidence?
- Is the information clear, accessible and provided at the correct time?

Do not use a blanket label

Not every AI-assisted email, graphic or edited document attracts the same obligation. Assess the legal category and context, then implement the right control.

Source note: European Commission transparency announcement and Article 50 [1], [9].

A 20-minute training video is not an operating model

Article 4 requires providers and deployers to take measures, to their best extent, to ensure a sufficient level of AI literacy among staff and other people dealing with AI systems on their behalf. The level should reflect their knowledge, experience, education, training and the context of use.

A proportionate programme might include:

Role	What competence should cover
Board and executive	Accountability, risk appetite, material decisions, escalation and evidence.
System owner	Purpose, limits, supplier position, performance, incidents and change.
Developer / technical team	Data, testing, security, logging, robustness and technical documentation.
Assessor / reviewer	Role, risk, prohibited practices, evidence quality and decision standards.
Front-line user	Approved use, limitations, disclosure, data handling and escalation.
Procurement / vendor manager	Contractual evidence, updates, incident notice and exit rights.

Keep the evidence

Record the required competence, training completed, date, role, assessment or confirmation and refresh cycle. Training is a control only if you can show who received what and why it was appropriate.

Source note: Article 4 and implementation timeline [2], [10].

Do not confuse regulatory classification with one generic score

A single 0-100 score cannot safely answer every governance question. Whether a practice is prohibited, whether an AI system is legally high-risk, whether the evidence is sufficient, and whether residual risk is acceptable are related but different decisions.

Illustrative AIRAS Cloud interface - synthetic data. Versioned dimensions, mandatory floors and a readable explanation trace.

A defensible method should separate:

- Legal applicability and role
- Prohibited-practice screening
- Inherent risk before controls
- Mandatory risk floors or deployment blockers
- Control requirements
- Evidence sufficiency and confidence
- Residual risk after verified controls
- The authorised human decision

High-risk dates

Annex III high-risk rules are scheduled for 2 December 2027. Annex I regulated-product rules are scheduled for 2 August 2028. Readiness work should start now, but future duties should not be presented as already applying in full.

A policy says what should happen. A control proves how it happens.

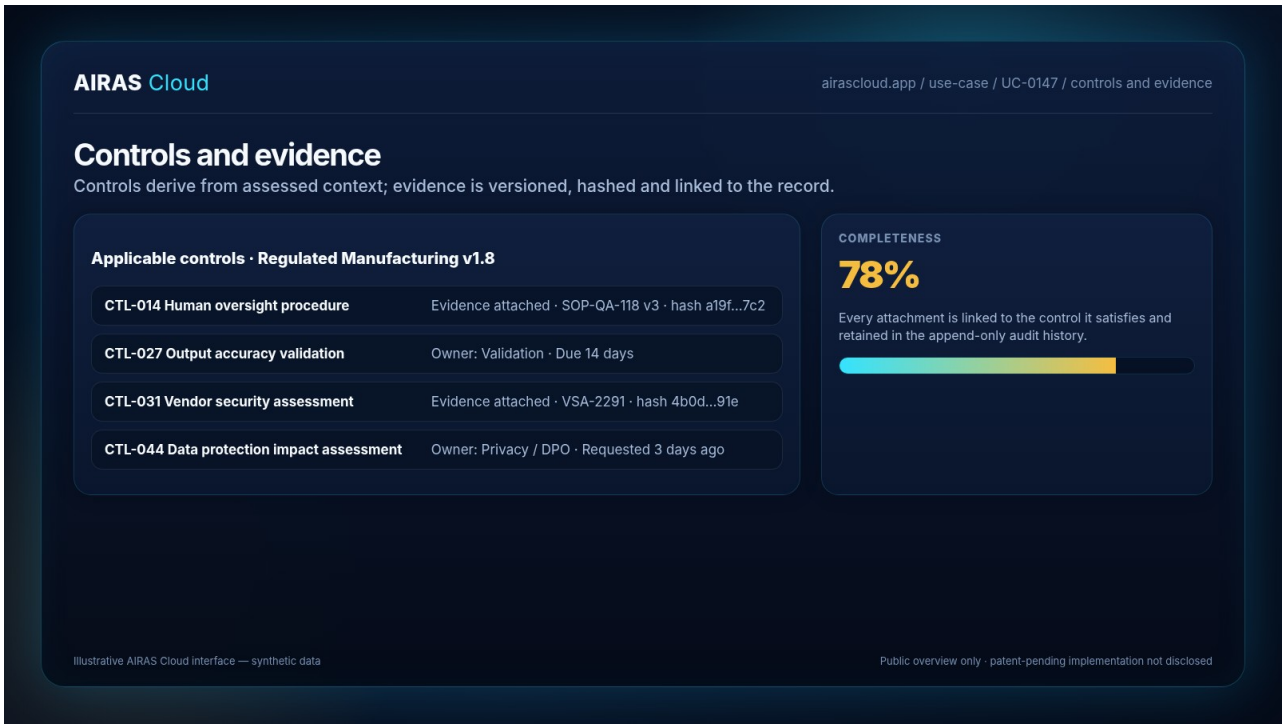

Illustrative AIRAS Cloud interface - synthetic data. Controls, owners, due dates and evidence completeness.

Controls should be derived from the actual context: what the system does, who it affects, what data it uses, where it can write back, how autonomous it is, what sector rules apply and what evidence is available.

Examples of control families

- Human oversight and intervention
- Data governance and quality
- Privacy and DPIA linkage
- Security, access and privileged actions
- Model and supplier change control
- Accuracy, robustness and performance
- Transparency and user information
- Logging, monitoring and incident response
- Validation and sector-specific assurance
- Exit, portability and business continuity

Control evidence is separate from control wording

A document saying a control exists does not prove it operates. Record the evidence, owner, review date, exceptions, test result and remediation status.

The decision needs a name beside it

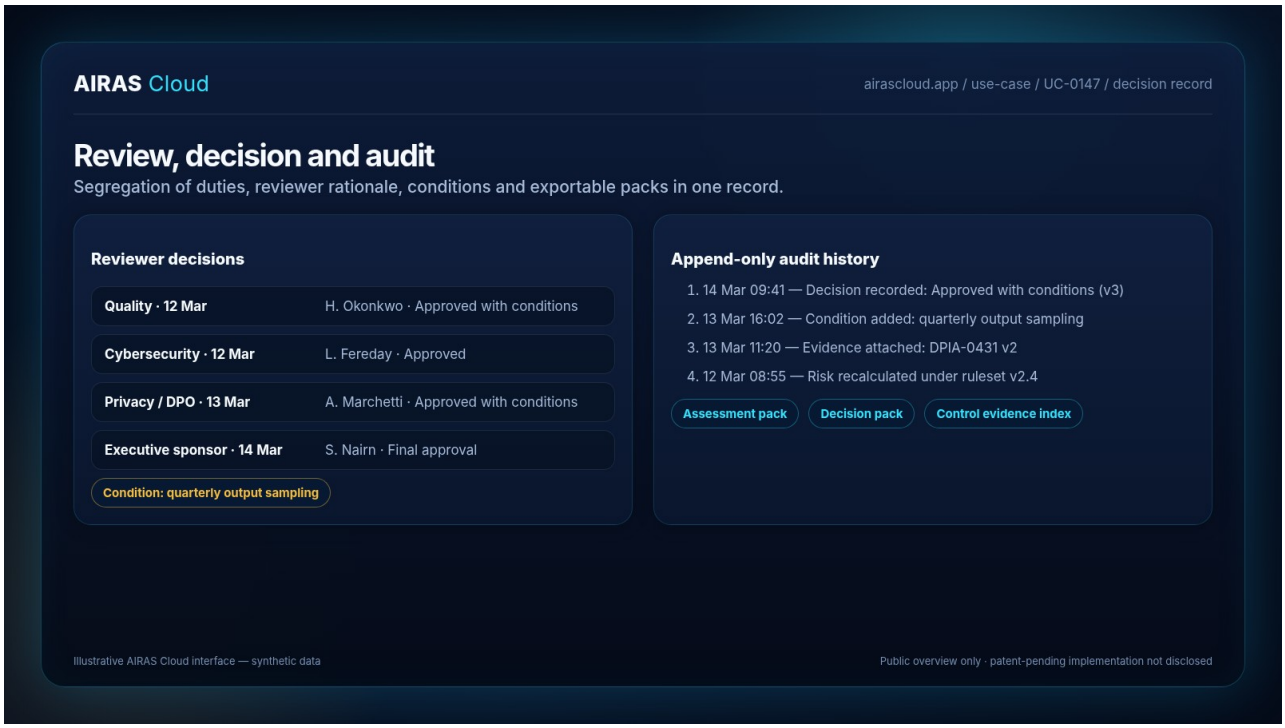

Illustrative AIRAS Cloud interface - synthetic data. Reviewer decisions, conditions and append-only history.

A good workflow separates the person submitting or assessing a use case from the person authorised to approve it where segregation of duties is required. It captures the final decision, the reasoning, any conditions and when reassessment is due.

The decision record should preserve:

- Input and evidence snapshot
- Ruleset and policy version
- Calculated outcome and mandatory rules
- Reviewers and their individual decisions
- Conditions, restrictions and expiry dates
- Final accountable decision
- Timestamped history that cannot be quietly overwritten
- Exportable assessment, decision and evidence packs

Evidence integrity is not evidence completeness

An append-only record can show that the history was not altered. It cannot prove that the original inputs were complete. Evidence sufficiency, uncertainty and human confirmation must remain visible.

Governance does not end at the approval meeting

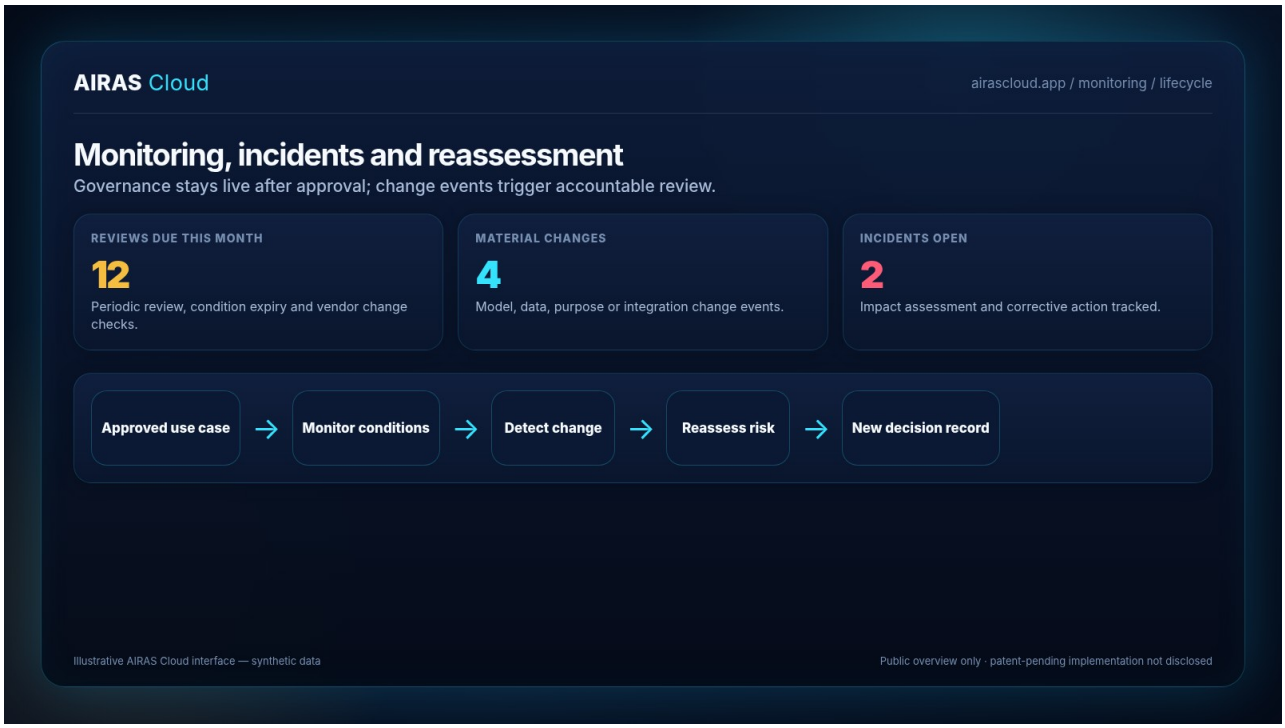

Illustrative AIRAS Cloud interface - synthetic data. Monitoring conditions, incidents and material-change reassessment.

AI systems change because models, suppliers, data, prompts, permissions, tools, users and business purposes change. A decision that was valid six months ago may no longer describe the system operating today.

Material-change triggers should include:

- Model or model-version change
- Supplier or contractual change
- New data source or data category
- Expanded user population or affected group
- New tool, API or write-back permission
- Wider or different business purpose
- Performance deterioration or drift
- Incident, complaint or override pattern
- Regulatory or policy change
- Control failure or expired evidence

Make monitoring actionable

A monitoring alert should identify the owner, required action, due date and whether the system can continue operating while the issue is resolved.

Vendor AI, shadow AI and agents

The systems most likely to be missed are not always the advanced models built by a data-science team. They are features switched on inside software the organisation already owns, tools adopted by staff, and agents that receive new permissions after an initial review.

Embedded vendor AI	Shadow AI	AI agents
Review release notes, feature flags, contracts, subprocessors, data use and the evidence the supplier can provide.	Use procurement, identity, browser, cloud, device, expense and staff-declaration signals. Record coverage and uncertainty.	Govern identity, tool permissions, budgets, approval gates, runtime actions, intervention and kill-switch capability.

Questions for every supplier

- What AI capability is present now, and what can be enabled later?
- What model and provider are used?
- What customer data is processed, retained or used for training?
- What documentation, logs and assurance evidence are available?
- How are model, policy and feature changes communicated?
- What incident, suspension, export and exit rights exist?
- Can the system be operated without creating an unacceptable concentration risk?

Commercial reality

A supplier saying “we are compliant” is not enough. Your organisation needs the evidence relevant to its own role, use, data, sector and controls.

The operating model is shared. The evidence changes by context.

Financial services & insurance

Model inventory, customer impact, outsourcing, explainability, credit or claims decisions, board risk acceptance and Central Bank expectations.

Healthcare & MedTech

Clinical workflow, patient safety, sensitive data, human oversight, validation, incident escalation and medical-device obligations.

Pharma & life sciences

GxP impact, computer-system validation, data integrity, supplier assurance, change control and inspection evidence.

Public sector

Prohibited-use screening, citizen impact, procurement, transparency, explainability, fundamental rights and public accountability.

Critical infrastructure

Operational technology, resilience, fallback, privileged access, write-back capability, cybersecurity and continuity.

Technology & SaaS

Embedded AI, customer role allocation, model-provider dependency, security, release control, data use and contractual evidence.

For SMEs

Proportionality matters. A smaller business does not need a large governance bureaucracy. It still needs a reliable inventory, named ownership, role-based literacy, a repeatable assessment, supplier evidence, human decisions and records proportionate to the risk.

The fine is not the only reason to act

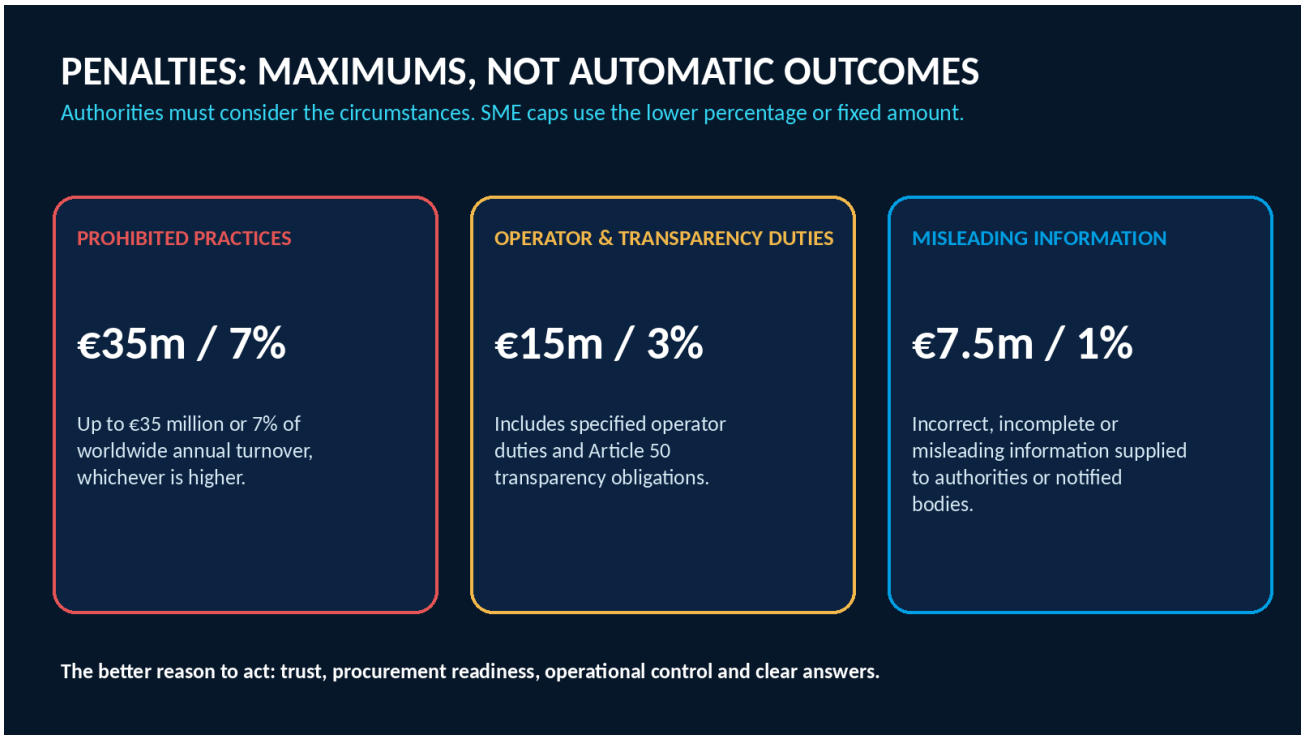

Statutory maximums under Article 99. For SMEs, including start-ups, the lower percentage or fixed amount applies.

Member States must provide effective, proportionate and dissuasive penalties. The maximum amounts are not automatic. Authorities consider factors including the nature, gravity and duration of the infringement, mitigation, cooperation and the size and interests of SMEs.

The commercial consequences can arrive first:

- Delayed procurement or failed customer assurance
- Audit findings and remediation cost
- Board or risk-committee escalation
- Restrictions on deployment or use
- Loss of confidence from customers, staff or the public
- Inability to answer a regulator or authority promptly
- Supplier disputes and unplanned exit work

The stronger business case

Governance is not just defensive. An organisation that can answer AI questions quickly can adopt useful AI faster, pass procurement more easily and give its board a clearer basis for decision.

Source note: Article 99 and European Commission transparency enforcement notice [1], [11].

A practical 90-day plan

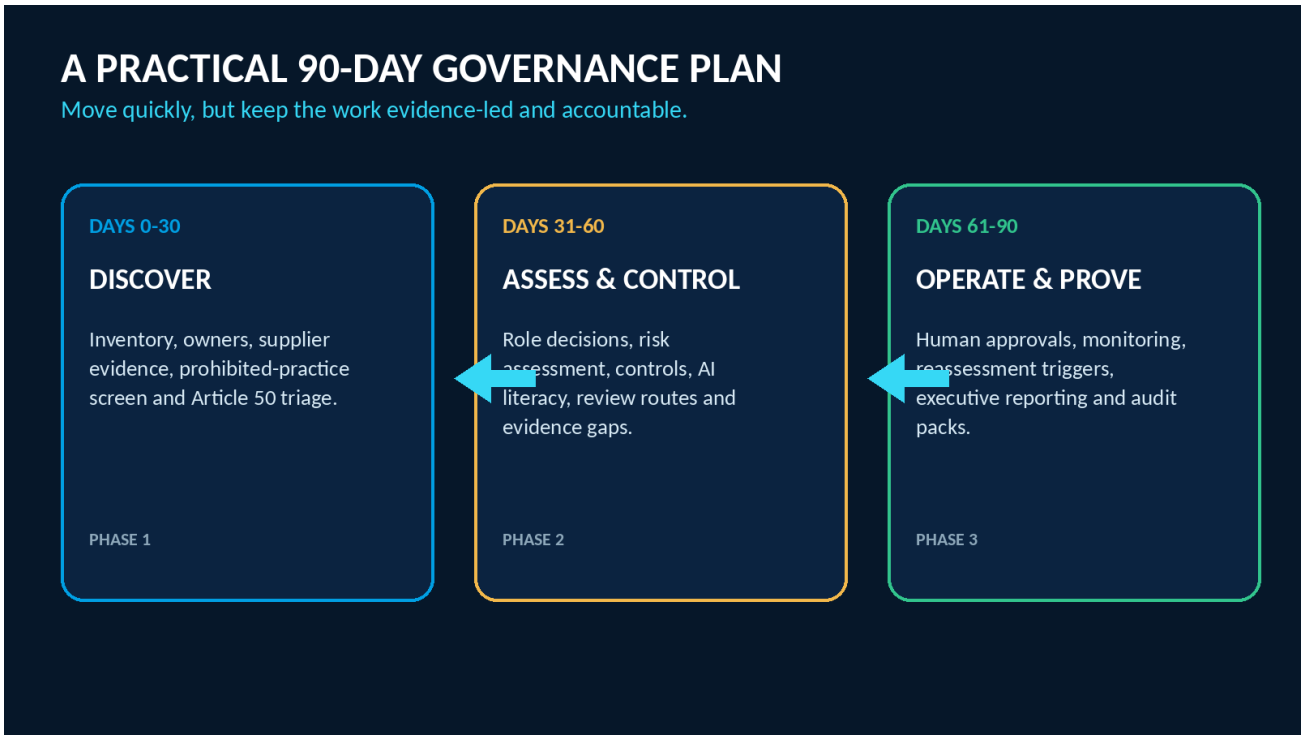

A proportionate implementation sequence for an organisation starting now.

Days 0-30: discover and stabilise

Appoint a sponsor. Build the inventory. Identify owners. Gather supplier evidence. Screen prohibited practices. Triage Article 50. Identify immediate data or security concerns.

Days 31-60: assess and control

Confirm roles. Apply a versioned assessment method. Map controls. Create role-based AI literacy. Establish review routes and close critical evidence gaps.

Days 61-90: operate and prove

Complete accountable decisions. Activate monitoring and reassessment. Produce board reporting. Test evidence packs. Integrate the process into procurement, change, security, privacy and risk.

Do not wait for perfection

Start with the highest-exposure systems and make uncertainty visible. A controlled, documented first version is better than an invisible perfect framework that never reaches operations.

How I turned the guidance into an operating platform

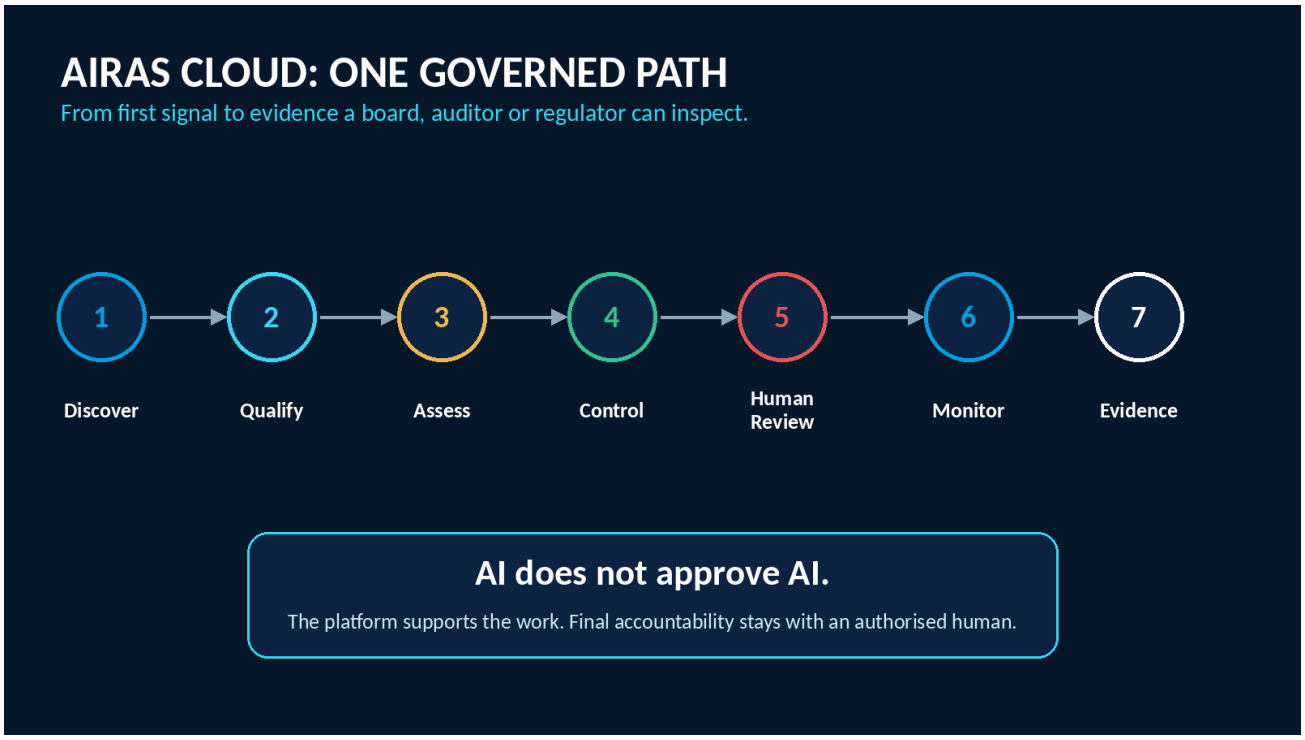

AIRAS Cloud brings the complete governance lifecycle into one governed record.

I built AIRAS Cloud because I kept seeing the same gap in regulated technology environments: organisations had policies, frameworks and committees, but the working evidence was spread across spreadsheets, emails, documents and individual memory.

AIRAS Cloud is designed to provide one controlled path from the first signal that AI may be present through qualification, assessment, controls, human review, monitoring and exportable evidence.

What AIRAS Cloud does not claim

It does not certify compliance, replace legal advice or guarantee that an organisation has disclosed every fact. It makes the process, evidence, uncertainty, rules and accountable decisions visible.

From intake to a decision you can defend

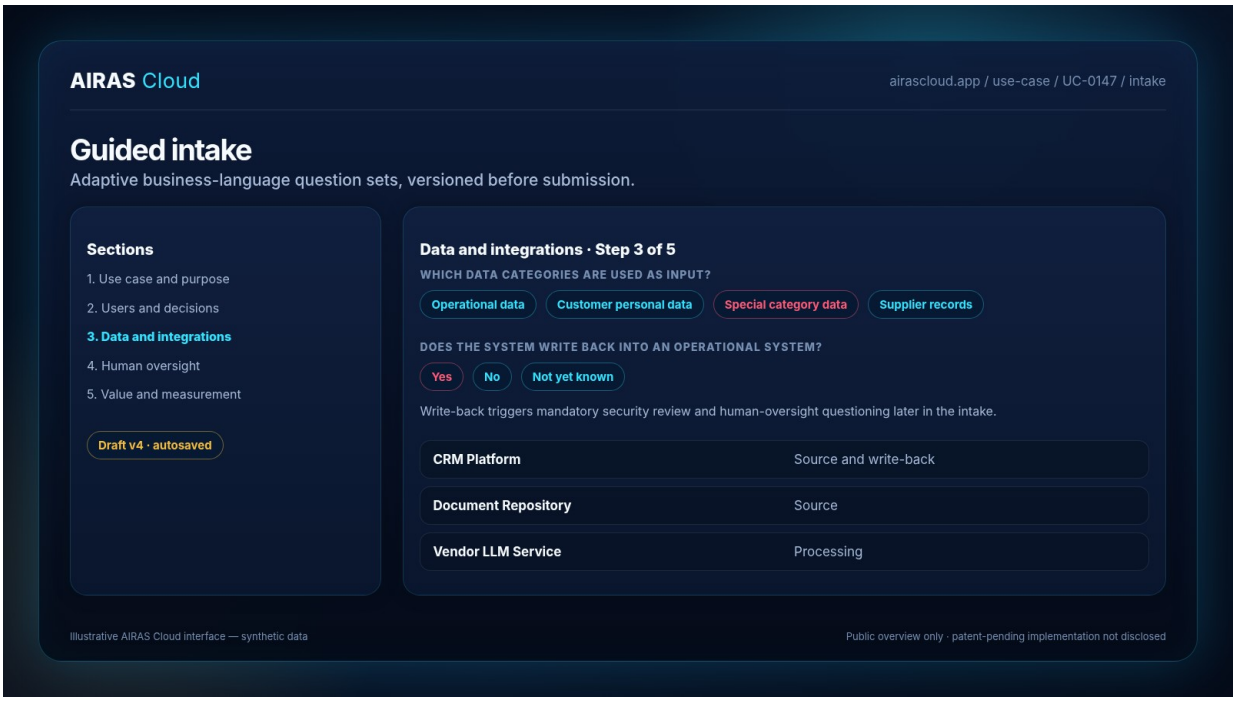

Guided intake - adaptive questions, evidence sources and visible status. Illustrative interface; synthetic data.

The intake records purpose, people, data, vendor, integrations, autonomy and evidence. It should ask only what the context needs, without hiding missing information behind a completed form.

Deterministic risk calculation - versioned rules, mandatory floors and explanation trace. Illustrative interface; synthetic data.

The assessment separates risk dimensions from mandatory rules and displays why the route changed. The same inputs and approved ruleset should reproduce the same result.

Controls, review, monitoring and evidence

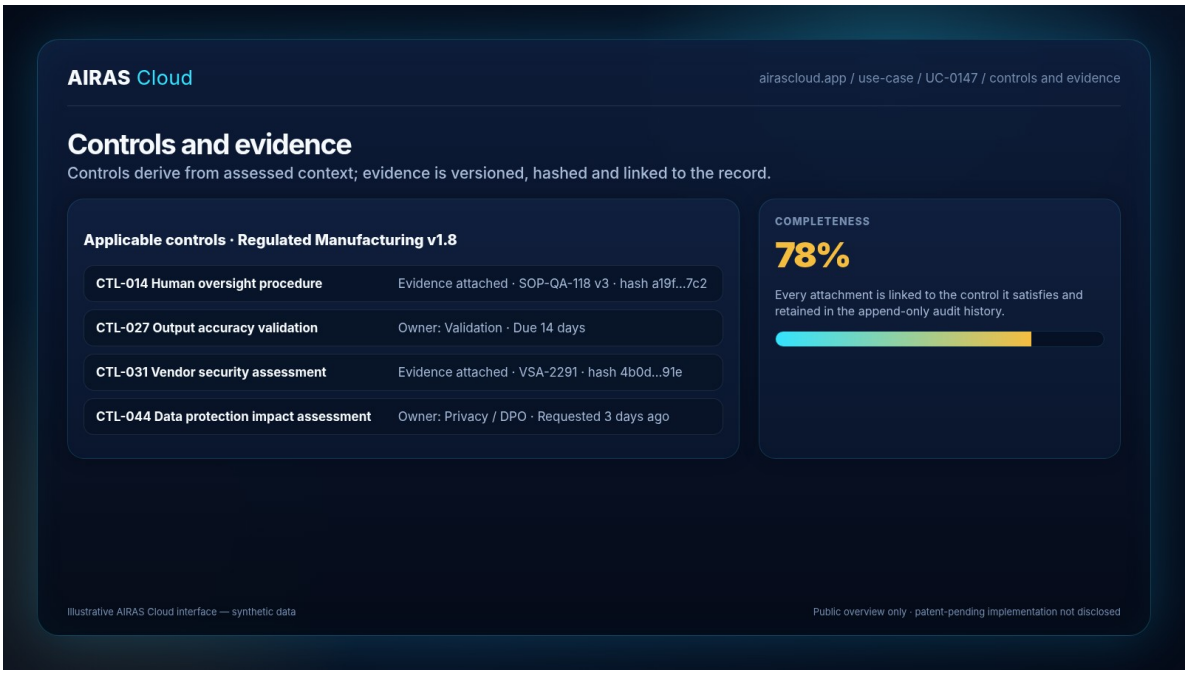

Controls and evidence - owners, due dates, versions and completeness. Illustrative interface; synthetic data.

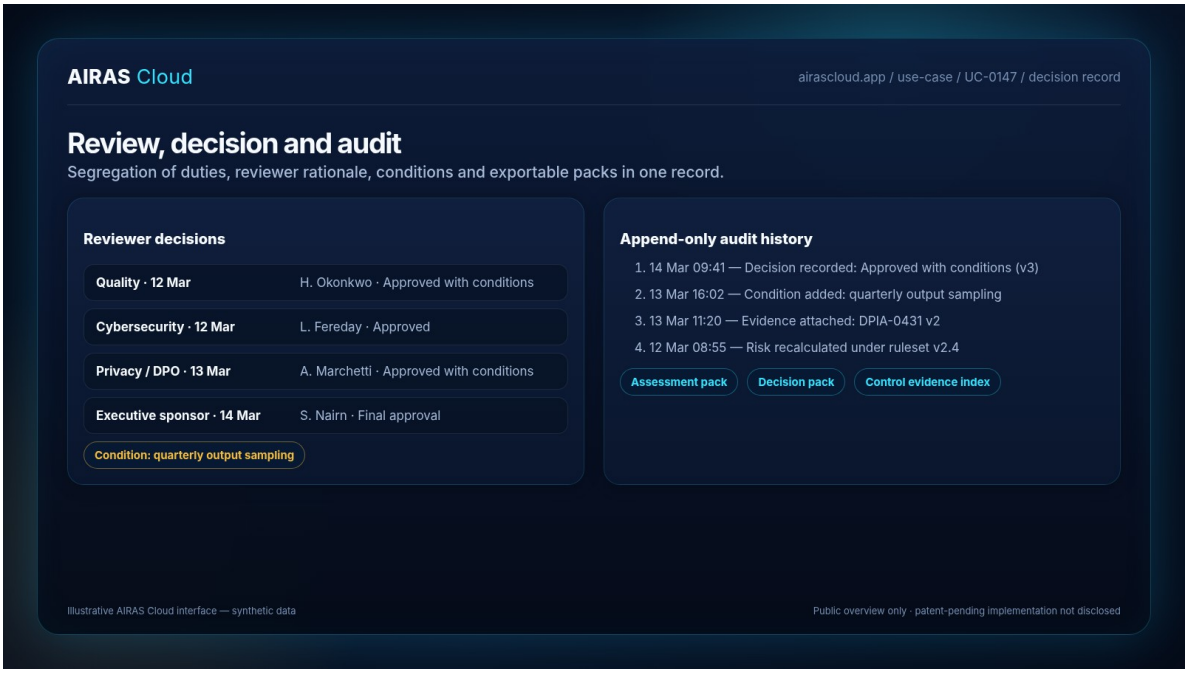

Reviewer decisions and append-only audit history. Illustrative interface; synthetic data.

The objective is not to make governance look busy. It is to create a record that an accountable person can use and that a board, auditor, customer or authority can reconstruct later.

Efficient, proportionate help for Irish organisations

Different organisations are at different stages. Some need to understand the Act. Others have a policy but no inventory. Others need to move a live portfolio into a controlled platform. AIRAS Cloud can be introduced in stages so the investment follows the actual need.

Executive briefing	A focused session for leadership covering the current legal position, business exposure, ownership model and immediate decisions.
Discovery and readiness sprint	A short engagement to reconcile the AI inventory, identify roles, screen current duties and produce a prioritised action plan.
Controlled enterprise pilot	An eight-week programme taking a representative portfolio through configuration, assessment, review, evidence and an integration proof.
Platform implementation	A scalable AIRAS Cloud environment with governed configuration, user roles, workflows, rulesets, monitoring and operational handover.

Cost-effective by design

The aim is not to sell a long consulting exercise. It is to establish the smallest credible operating model that matches the organisation's AI estate, regulatory exposure and evidence needs - then scale it as adoption grows.

Contact: interest@airascloud.com

Website: airascloud.com

Can your organisation answer these questions today?

- ☐ Can we produce a current inventory of AI systems, models, agents and embedded vendor features?
- ☐ Does every material system have a named business owner?
- ☐ Have we determined our role for each system?
- ☐ Have prohibited practices been screened and hard stops enforced?
- ☐ Have Article 50 transparency duties been assessed?
- ☐ Can we evidence proportionate AI literacy by role?
- ☐ Is there one approved, version-controlled assessment method?
- ☐ Are mandatory issues prevented from being averaged away?
- ☐ Is the assessor separated from the approver where required?
- ☐ Can we show the evidence, assumptions, rules and decision?
- ☐ Are control owners, due dates and evidence visible?
- ☐ Do we monitor incidents, conditions and material change?
- ☐ Can a customer, auditor or authority receive a coherent evidence pack?
- ☐ Can we export our data and continue operating if a supplier fails?

A simple interpretation

If several answers are “no”, the priority is not another policy. It is a controlled discovery and implementation plan.

Five points worth repeating

1	**Key message** 2 August 2026 is a major enforcement milestone, but it did not make every AI system high-risk overnight.
2	**Key message** The first governance question is not “Do we have a policy?” It is “What AI are we actually using?”
3	**Key message** Irish businesses need to assess their role and obligations per AI system, not apply one label to the whole company.
4	**Key message** Transparency, AI literacy, prohibited-practice screening, ownership and evidence are live operating issues now.
5	**Key message** AI can support governance work. The final accountable approval must remain human.

Suggested attribution

Richie Higgins, Founder & CEO of AIRAS Cloud

“The EU AI Act is not asking Irish businesses to stop using AI. It is asking them to know what they are using, control it properly and be able to prove the decisions they made.”

Media contact: interest@airascloud.com | airascloud.com

Clear answers for Irish organisations

Does the EU AI Act apply to Irish businesses?

Yes. The Regulation applies directly in Ireland, subject to its scope, exemptions, role definitions and staged implementation dates.

Did every AI system become high-risk on 2 August 2026?

No. Article 50 transparency and enforcement of currently applicable rules are central to the milestone. The amended high-risk dates are later.

Does every AI-generated item need a label?

No. Article 50 obligations depend on the system, output and context. Organisations should assess the specific category rather than use one blanket label.

What should an SME do first?

Create a proportionate inventory, name owners, determine roles, screen prohibited practices, assess transparency and establish a repeatable human-governed process.

Can a supplier say it is compliant and close the issue?

No. Supplier evidence matters, but your organisation must still assess its own role, use, data, controls and responsibilities.

Does AIRAS Cloud guarantee compliance?

No. AIRAS Cloud supports governance workflow, structured assessment, controls, human review, monitoring and evidence. It does not replace legal advice or remove accountability.

Official sources used for the regulatory position

[1] European Commission, “Safer and more transparent AI”, 2 August 2026. https://commission.europa.eu/news-and-media/news/safer-and-more-transparent-ai-2026-08-02_en

[2] European Commission AI Act Service Desk, implementation timeline. https://ai-act-service-desk.ec.europa.eu/en/ai-act/eu-ai-act-implementation-timeline

[3] Regulation (EU) 2026/1744, Digital Omnibus on AI. https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng

[4] Regulation of Artificial Intelligence Act 2026, Act No. 31 of 2026. https://www.irishstatutebook.ie/eli/2026/act/31/enacted/en/pdf

[5] Houses of the Oireachtas, Regulation of Artificial Intelligence Act 2026. https://www.oireachtas.ie/en/bills/bill/2026/69/

[6] Department of Enterprise, Tourism and Employment, “The EU AI Act and my organisation”. https://enterprise.gov.ie/en/what-we-do/innovation-research-development/artificial-intelligence/eu-ai-act/eu-ai-act-and-my-organisation.html

[7] European Commission guidelines on the definition of an AI system. https://ai-act-service-desk.ec.europa.eu/sites/default/files/2026-01/guide-definition_en.pdf

[8] EU AI Act Service Desk, Article 5: prohibited AI practices. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-5

[9] EU AI Act Service Desk, Article 50: transparency obligations. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-50

[10] EU AI Act Service Desk, Article 4: AI literacy. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-4

[11] EU AI Act Service Desk, Article 99: penalties. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-99

[12] Regulation (EU) 2024/1689, Artificial Intelligence Act. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

Legal qualification

This guide provides general information for business planning. Obligations depend on the organisation’s role, the system, intended use, sector and facts. Where this guide and the legal text differ, the legal text governs.

The pressure is real. The response can still be controlled.

Start with the inventory. Establish ownership. Apply one defensible process. Keep the evidence.

AIRAS Cloud brings that work into one governed operating layer.

airascloud.com | interest@airascloud.com

AIRAS Cloud - AI Governance Made Operational

AI does not approve AI.