



SHADOW AI DEFENCE / A CEO PERSPECTIVE

Solve the AI problem you didn't know you had. Before it costs you.

AI can arrive through the software you already trust.
AIRAS brings the evidence, assessment and accountability together.

RICHARD HIGGINS

Founder and Chief Executive Officer

AIRAS Cloud | 24 October 2026

A PERSPECTIVE FROM OUR FOUNDER

Solve the AI problem you didn't know you had. Before it costs you.

Why AIRAS Cloud exists: to uncover overlooked AI use and turn uncertainty into a decision your organisation can defend.

RICHARD HIGGINS

Founder and Chief Executive Officer | AIRAS Cloud

24 October 2026



Imagine signing off a successful software enhancement. Your recruitment team has the candidate-matching feature it wanted, testing is complete and an authorised administrator has enabled it. Hiring managers can review applications more efficiently, using a platform the organisation has trusted for years. Recruiters still make the final hiring decision.

The supplier was already approved. The agreement was already in place. Your team checked integration, access, functionality, cost and usability. There was no new vendor to investigate and no unauthorised website to block.

Then someone asks what AI the feature uses, which applicant information it processes and how strongly its recommendations influence who gets considered. What have candidates been told? Who assessed this particular use? Where is that decision?

In this illustrative organisation, those questions were never brought together. The enhancement solved the operational problem, but its AI implications remained between procurement, HR, technology and risk. For a candidate waiting to hear about an interview, the difference between a helpful feature and an unexamined influence on selection is hardly academic.

The technology is real. Greenhouse publicly documents AI-assisted Talent Matching, administrator-controlled activation and the fact that the feature does not automatically advance or reject candidates.^{[1][2]} This is not a reported incident involving Greenhouse or its customers. It shows how a legitimate, disclosed capability can still need a separate governance decision inside the organisation using it.

You approved the supplier. Who assessed the AI?

That question captures why AIRAS Cloud exists. I want organisations to discover these gaps while they still have time to make an informed decision, with a clear process for establishing the facts and acting on them.

The gap inside an ordinary change

What concerns me is how easily this can happen when capable people are doing their jobs properly. The business wants a solution. Technology teams need a working integration. Procurement needs an appropriate agreement. Those responsibilities are substantial, and a technical description does not automatically become an AI assessment simply because somebody has read it.

Supplier approval reflects the service, intended use and evidence considered at the time of review. A later enhancement may introduce a model, a processing arrangement or an influence on decisions that the original approval never covered. The supplier may explain the change accurately; the organisation still needs to connect that explanation to its own use, responsibilities and controls.

This is one form of shadow AI: AI operating outside established governance, including capabilities within approved software. Processing may happen in the supplier's infrastructure or through its backend providers while staff remain inside the permitted application.

A browser policy, DNS filter, secure web gateway or firewall rule can restrict identified destinations. A blocking rule cannot, by itself, establish a complete AI inventory, determine legal classification or record an accountable decision about use. Its presence does not establish what an approved cloud service does with information.

We built AIRAS Cloud to reduce dependence on someone happening to notice the relevant detail and knowing which questions to ask next. Our onboarding gives that detail a route into evidence, verification, assessment and ownership.

What finding out late can cost

Discovering unassessed AI can create an investigation that nobody planned to resource. The organisation must establish what happened, who was affected, what the supplier can substantiate and whether the feature should remain enabled. Decisions or outputs may need review, potentially after information has already been processed.

That work can consume the time of legal advisers, privacy specialists, technology leads and operational managers. It competes directly with delivery. Those hours have a cost, even when no regulator is involved.

Consider a separate, hypothetical services business with €20 million of annual revenue, €14 million of cost of sales and €4 million of other operating expenses. Its gross margin is 30% and operating profit is €2 million.

Assume unassessed AI in customer delivery creates €300,000 of additional direct delivery costs: €180,000 to reperform work, €80,000 for temporary delivery capacity and €40,000 to verify corrected deliverables. Another €200,000 of operating expenses covers legal and privacy advice, technical investigation, supplier review and corrective training.

Annual financial measure	Original position	After the assumed costs
Revenue	€20,000,000	€20,000,000
Cost of sales	€14,000,000	€14,300,000
Gross profit	€6,000,000	€5,700,000
Gross margin	30.0%	28.5%
Other operating expenses	€4,000,000	€4,200,000
Operating profit	€2,000,000	€1,500,000
Operating margin	10.0%	7.5%

Gross margin falls by 1.5 percentage points, gross profit by €300,000 and operating profit by 25%. At the original 30% gross margin, replacing that €300,000 of gross profit requires €1 million of additional sales, assuming the same margin.

These assumptions illustrate the mechanics; they are not typical incident costs, a customer result or a claimed AIRAS saving. Only direct delivery costs reduce gross profit in this model. The other response costs reduce operating profit. An internal recruitment review would often affect operating expenses instead, with classification depending on the work and accounting policy. Fines, lost revenue, compensation and insurance recoveries are excluded.

For a business working hard to earn its margin, corrective work on that scale deserves attention before a regulator becomes involved.

The legal position must be understood

In recruitment, Annex III of the EU AI Act includes certain systems intended to analyse or filter applications and evaluate candidates, subject to Article 6's classification rules. A human retaining the final hiring decision does not, on its own, settle classification. Purpose, function and influence still need examination.^{[3][4]}

Article 99 provides substantial administrative fine ceilings for defined infringements:^[5]

Infringement category	Headline maximum
Prohibited AI practices under Article 5	€35 million or 7% of worldwide annual turnover for the preceding financial year
Specified operator and transparency obligations	€15 million or 3%
Incorrect, incomplete or misleading replies to relevant authorities or notified bodies	€7.5 million or 1%

For undertakings, the higher amount generally applies. SMEs, including start-ups, benefit from the lower comparison; amended rules also apply a lower comparison to small mid-cap companies for the latter two categories. Percentages refer to worldwide annual turnover for the preceding financial year. These are statutory ceilings considered against the actual infringement and circumstances, not automatic charges for finding an unregistered feature.

For the more serious categories of GDPR infringement, penalties can reach €20 million or 4% of worldwide annual turnover for the preceding financial year, whichever is higher for an undertaking. Authorities can also limit or prohibit processing. These maximums and AI Act maximums must not be added together as an inevitable bill.^[6]

The relevant high-risk provisions for Annex III systems apply from **2 December 2027**, subject to transitional rules. Other provisions already apply. There is no universal deadline requiring reassessment of every historical project; applicability depends on the use, role and obligations involved.^{[7][8]} The practical response is to establish those facts early enough to act.

How AIRAS makes the question manageable

AIRAS Shadow AI Defence starts with evidence the organisation authorises for review: vendor material, contracts, technical specifications, architecture records and project documents. Analysis identifies candidate AI capabilities, with source passages and confidence indicators so reviewers can see why a finding was raised and confirm or dismiss it with a recorded reason.^[9]

In our recruitment example, the candidate-matching description provides a specific question for the owner and supplier. Is the feature included, enabled and used? Which workflows depend on it? What information does it process?

Human confirmation matters here. A product offering AI is not proof of customer activation. A marketing reference, an unused option and a deployed feature require different conclusions. Missing information remains a question to resolve.

Once the facts are established, ARIE, AIRAS Cloud's regulatory and risk assessment engine, applies versioned, deterministic logic to the recorded context. It supports AI qualification, role identification, prohibited-practice screening, classification and obligation mapping. Its reasoning is traceable to the inputs, legal sources and ruleset. A prohibited use is a gate that a favourable score cannot override.^[10]

AIRAS connects the use case to owners, controls, evidence requirements and role-based review. Decisions and conditions sit in the same governed record, ready for reassessment when things change.^[11] A named lead coordinates specialist input from privacy, security, procurement and the business owner.

The practical value is having a defined decision and shared evidence to work from, reducing the need to assemble the process afresh. Specialist expertise remains essential, and complex cases require substantial review. AIRAS makes that work focused and repeatable.

The outcome can be approval with controls, a request for evidence, a restriction or a decision not to proceed. The record captures who decided, why, the conditions and who must implement them.

A decision people can stand behind

AIRAS explains its scores through the dimensions, rules and reasoning applied.^[12] Discovery confidence, legal classification and operational exposure after controls answer different questions. High confidence is not high risk. A low score cannot override prohibition, and missing evidence cannot establish low risk. I would only score the recruitment example after verifying its deployment facts and controls.

People retain the authority to approve use and the responsibility to oversee it. That accountability is central to AIRAS; every assessment depends on the scope of review and the evidence available.

Take one recently enabled feature in a platform you trust. Can you identify its AI assessment, accountable owner and approval record? If that answer depends on somebody's memory, AIRAS can help bring the evidence and decisions together.

AIRAS Cloud gives organisations a practical way to bring overlooked AI under accountable control. That is why we do this work. My invitation is to bring us one existing supplier, one recently enabled feature or one proposed enhancement. We can show how the process works, while you still have room to decide what happens next.

Richard Higgins
RICHARD HIGGINS

Founder and Chief Executive Officer, AIRAS Cloud

airascloud.com | interest@airascloud.com



Sources and publication notes

The organisational scenarios and financial figures are illustrative. Legal references reflect sources reviewed on 24 October 2026; applicability depends on the system, the organisation's role and relevant transitional provisions.

1 Greenhouse Support, "Talent Matching".

<https://support.greenhouse.io/hc/en-us/articles/41396009937307-Talent-Matching>

2 Greenhouse Support, "Talent Matching setup".

<https://support.greenhouse.io/hc/en-us/articles/45126074890523-Talent-Matching-setup>

3 European Commission, AI Act Service Desk, Annex III.

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3>

4 European Commission, AI Act Service Desk, Article 6.

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-6>

5 European Commission, AI Act Service Desk, Article 99.

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-99>

6 European Data Protection Board, "Data protection authority and you".

https://www.edpb.europa.eu/sme/find-practical-info/data-protection-authority-you_en

7 European Commission, AI Act Service Desk, Article 113.

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-113>

8 European Commission, AI Act Service Desk, Article 111.

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-111>

9 AIRAS Cloud, "Shadow AI: finding the AI you have not governed".

<https://airascloud.com/resources/shadow-ai>

10 AIRAS Cloud, "ARIE".

<https://airascloud.com/arie>

11 AIRAS Cloud, "How AI Governance Works".

<https://airascloud.com/how-it-works>

Author and brand imagery

Published author portrait: [AIRAS Cloud public media resources](#). Logo supplied by AIRAS Cloud. The cover combines the original AIRAS identity with bespoke abstract brand artwork. The badge denotes 2026 Best Company Awards finalist recognition.