AIMONGER
WHITE PAPER

AI Discovery

EU AI Act Board Readiness: A Practical Checklist for 2026

Regulation (EU) 2024/1689 entered into force in August 2024 with obligations that phase in through 2027. Most mid-market firms are not training frontier models; they are deploying vendor AI in hiring, credit, customer service, and document workflows. The board variable is whether the company knows what it runs, who owns it, and whether the stack was built for audit or for demo speed.

David Saliba . 2026-06-27 . aimonger.com/whitepapers/eu-ai-act-board-readiness/

The problem in one sentence

The EU AI Act is no longer a future briefing note. It is a staggered compliance calendar running in parallel with a sharp rise in enterprise AI deployment, and most firms cannot yet say what they run, who owns it, or where the human gates sit.

Walk into a steering meeting and someone will say “we are not an AI company.” Walk into procurement, HR, and customer operations of the same firm and you often find something different: résumé screening features in an ATS, fraud scoring in payments, document extraction in finance, copilots in support, and chat widgets on the website. Some of those are low-risk productivity tools. Some touch Annex III categories. Almost none sit in a maintained register with a named owner and an incident path.

That gap is the 2026 board problem. Legal text is necessary but not sufficient. Readiness is an operating artefact: a register, gates, documentation, vendor diligence, and a cadence that survives staff turnover. The firms that already run agentic workflows with logging and human approval are closer than firms that deployed a chatbot with no audit trail and called it innovation.

> **EU AI Act readiness, plainly:** a documented position on what AI the company deploys, how each system is classified, who owns it, and what evidence exists when a regulator, customer, or auditor asks, not a one-off legal opinion filed and forgotten. **In practice:** a mid-market insurer maintains a living register of 14 AI-enabled workflows, each with purpose, data sources, risk tier, human gate, and vendor contract clause; the CIO can answer board question three in five minutes without opening a ticket queue.

> **High-risk (Annex III), in short:** uses listed in the regulation where flawed or biased outputs can materially affect access to employment, credit, insurance, education, or similar rights, triggering stronger documentation, oversight, and monitoring duties for deployers when in scope. **Worked case:** if your hiring platform ranks candidates automatically and nobody can explain the feature set or produce evaluation records, you are in board conversation territory even if nobody used the words “high-risk” in the project charter.

Why this landed now

Three forces converged in 2025 and 2026.

First, the Act's clock started. Prohibited practices became enforceable from February 2025. Obligations on providers of general-purpose AI models applied from August 2025. Transparency duties for certain AI interactions arrive in August 2026. High-risk system requirements for many deployers consolidate through August 2027. "We will deal with it next year" is no longer a calendar that finance and HR can hide behind.

Second, deployment outran governance. Eurostat reports 20.0 per cent of EU enterprises with 10 or more employees used at least one AI technology in 2025, up from 13.5 per cent in 2024. That is a 6.5 percentage point jump in one year. Vendor SaaS shipped AI features on default settings. Teams adopted them because the button was there, not because anyone classified risk.

Third, accountability expectations shifted. McKinsey's State of AI 2025 survey finds nearly nine in ten respondents say their organisations regularly use AI in at least one function, while only 39 per cent report any enterprise-level EBIT impact and about 6 per cent qualify as high performers attributing more than 5 per cent of EBIT to AI. Boards approved spend in the access phase. They will ask for proof and controls in the accountability phase. The EU AI Act is one face of that broader demand.

This paper is a practical operating briefing for executives and CIOs. It is not legal advice. It translates the board's job into eight questions, evidence you should expect to see, and decision criteria investment committees can use without reading the full regulation.

Eight questions for this quarter's board meeting

Each question below is the headline. The prose is what "answered" looks like in an mid-market firm that serves EU customers.

1. Inventory: What AI systems are we running today?

Board reading: If the answer is "we will ask IT," the firm is not ready. AI is embedded in SaaS, not only in internal data-science projects.

An inventory is a maintained register, not a one-off survey. Minimum fields per system:

- **Name and vendor** (including AI features inside existing CRM, ATS, ERP, or support platforms)
- **Purpose** (what decision or workflow it affects)
- **Data sources** (what it reads at inference time)
- **Risk tier** (prohibited / high / limited / minimal, with rationale)
- **Owner** (business name, not "IT")
- **Human gate** (where a person must approve, override, or review)
- **Documentation location** (technical file, model card, vendor attestation)

Include shadow paths. If staff paste customer data into consumer chat tools because the approved tool is too slow, the inventory is incomplete and the risk register is wrong.

> **AI system register means:** a single list of every AI-enabled workflow and vendor feature the company relies on, updated when procurement signs, HR turns on a module, or engineering ships. **On Monday:** when finance enables invoice extraction in the ERP, procurement adds the vendor's AI addendum, and the register row appears before go-live, not after audit discovery.

2. Risk tier: Which systems touch Annex III high-risk categories?

Board reading: Classification is not a branding exercise. It follows use case and scope rules in the regulation.

High-risk lists in Annex III cover domains such as employment and worker management, access to essential services, creditworthiness, certain insurance activities, education, and law enforcement, among others. Exceptions exist; legal counsel must confirm. Your operating job is to flag candidates early.

Run a structured pass:

1. **Map workflows to domains.** Hiring, promotion, scheduling, and performance tools often touch employment. Credit and collections touch scoring. Underwriting engines touch insurance pricing.
2. **Separate provider vs deployer duties.** You may be a deployer of a vendor's high-risk feature even if you did not train the model.
3. **Document the rationale.** "Minimal risk" without reasoning is not a position. "High-risk hiring ranking; human review before interview shortlist; vendor contract section 14" is a position.

3. Data lineage: Where does training and inference data come from? Does any leave the EU?

Board reading: Regulators and enterprise customers increasingly ask where data flows, not only which model brand you use.

For each inventoried system, trace:

- **Training data** (if you fine-tune or vendor claims domain adaptation)
- **Inference inputs** (documents, transcripts, biometrics, behavioural logs)
- **Retention** (logs, prompts, outputs, embeddings)
- **Geography** (region of processing and subprocessors)
- **Lawful basis** (contract, consent, legitimate interest where relevant under GDPR)

Mid-market firms often discover that a US-hosted SaaS AI feature processes EU personal data by default. That is not automatically unlawful, but it must be documented, contracted, and aligned with transfer tools and customer commitments.

4. Human oversight: Which decisions have a named human gate, and which run fully automated?

Board reading: "Human in the loop" on a slide is not oversight. Oversight is a named role, a defined trigger, and evidence.

For each consequential workflow, specify:

- **What the system may recommend vs decide**
- **Who may override** (job title and delegation)
- **When escalation is mandatory** (edge cases, low confidence, protected characteristics proximity)
- **What is logged** (input summary, output, reviewer identity, override reason)

Agentic architectures that log tool calls, enforce policy at an orchestration layer, and block dispatch without approval are structurally easier to audit than monolithic black-box APIs with no trace. The board question is whether your stack was built for accountability or for demo speed.

5. Transparency: Do customers and employees know when they interact with AI?

Board reading: Transparency obligations phase in for certain interactions from August 2026. Customer trust often demands earlier disclosure.

Checklist per channel:

- **Support chat and voice bots** (clear disclosure that AI is involved; path to human)
- **Generated content** (marketing, reports, outbound correspondence where AI drafted material portions)
- **Employee-facing tools** (monitoring, scheduling, performance analytics)
- **Recruitment** (automated parsing, ranking, or video analysis)

Plain language beats buried footnotes. If a reasonable employee or customer would not know AI shaped the interaction, fix the UX before legal asks.

6. Documentation: Do we have technical documentation for each deployed system?

Board reading: Documentation is how you prove diligence months later when the original engineer has left.

Expected artefacts vary by tier. For high-risk candidates, plan for:

- **System description and intended purpose**
- **Data governance and quality measures**
- **Risk management measures and test results**
- **Human oversight design**
- **Accuracy, robustness, and cybersecurity measures**
- **Change log** (model version, prompt, retrieval corpus updates)

For vendor SaaS, much of this sits in supplier documentation. Your job is to obtain, index, and gap-fill where the vendor file is thin. A folder per system beats a shared drive of unsigned PDFs.

7. Incident process: Who gets notified if an AI system produces a harmful or biased output?

Board reading: Incidents will happen. Unowned incidents become crises.

Define:

- **What counts as an AI incident** (wrongful denial, discriminatory outcome, data leak via prompt, runaway agent action)
- **First responder** (service owner, not anonymous inbox)
- **Escalation path** (legal, DPO, communications, board notification threshold)
- **Containment** (disable feature flag, rollback model version, freeze agent tools)
- **Post-incident review** (root cause, register update, customer notification decision)

Run one tabletop exercise before you need it. The register owner and on-call name should appear in the board pack.

8. Vendor contracts: Do our AI SaaS providers carry EU AI Act obligations in their terms?

Board reading: You cannot contract away all deployer duties, but you can refuse vendors who offer no compliance artefacts.

Procurement should require, for AI features:

- **Role clarity** (provider vs deployer vs importer)
- **Documentation delivery** (model cards, system documentation updates)
- **Incident cooperation**

- **Subprocessor transparency**
- **Change notification** (model swaps, retraining, new data use)
- **Audit and inspection rights** where appropriate

Renegotiate renewals on the calendar, not after a failed audit. The firms that waited until August 2027 to read SaaS terms will pay rush pricing.

What "ready" looks like for a mid-market company

Readiness is not a certification badge. It is a documented position the executive team can restate quarterly:

1. **Living AI system register** with owners and risk tiers, updated on procurement events
2. **Policy gates** on high-risk workflows (hiring, credit, insurance, biometric identification)
3. **Staff training** on limitations, refusal paths, and escalation (operators, not town-hall spectators)
4. **Vendor due diligence** on any AI feature embedded in existing SaaS
5. **Human oversight design** with logs that survive audit
6. **Incident playbooks** with named roles
7. **Legal sign-off** on classification edge cases, recorded in writing

Companies that already run governed agentic workflows with evaluation sets and rollback are closer than companies that deployed a public chatbot with no logging. Compliance favours architectures that were designed to be inspected.

Evidence base: adoption scale vs governance depth

Evidence	Finding	What it does and does not prove
Eurostat enterprise survey (2025)	20.0% of EU enterprises with 10+ workers used at least one AI technology, up from 13.5% in 2024	Official adoption scale; does not measure compliance readiness or depth of use
EU AI Act (Regulation 2024/1689)	Staggered obligations from 2025 through 2027 by role and system class	Legal framework; national enforcement details still evolving
McKinsey State of AI (2025)	88% report regular AI use in at least one function; 39% report any enterprise EBIT impact; ~6% high performers (>5% EBIT)	Global executive survey; shows deployment ahead of measured control and value
EC AI Office / Commission guidance	Provider and deployer duties, GPAI model documentation, prohibited practices	Official interpretation; must be paired with counsel on specific systems
Deloitte State of AI in the Enterprise (2026)	25% of surveyed leaders moved 40%+ of AI experiments to production; 30% redesigned key processes	Suggests many systems reach production without process or documentation maturity

The conclusion is not that every firm is non-compliant. It is that adoption scale and governance depth are on different curves. A board that tracks licence counts but not register completeness will be surprised by the compliance curve.

Operating implications for CIO, legal, HR, and procurement

Compliance is a matrix, not a department.

CIO and engineering own inventory truth, logging, human-gate implementation, and change control. If the register lies, every other function optimises fiction.

Legal and DPO own classification, contract language, transfer analysis, and regulatory interface. They cannot classify systems that engineering never disclosed.

HR owns employment-related AI features end to end. Hiring tools are the most common Annex III surprise in mid-market firms.

Procurement owns vendor evidence at renewal. Default AI toggles in SaaS are procurement events even when no new PO was signed.

The board owns calendar and escalation. It sets the expectation that unanswered questions from the list of eight are treated like an open control deficiency, not an academic exercise.

Suggested ninety-day cadence:

1. **Weeks 1-2:** Draft register from IT asset lists, procurement renewals, and workflow interviews; flag shadow IT.
2. **Weeks 3-4:** Risk-tier pass with legal on top ten systems by business impact; assign owners.
3. **Weeks 5-8:** Close human-gate and logging gaps on high-impact workflows; begin vendor documentation chase.
4. **Weeks 9-12:** Tabletop incident; board readout with open gaps, dates, and budget for remediation.

Cross-read the adoption picture in State of AI adoption in Malta and European business and the workflow value gap in Absorption capacity. Compliance without absorption produces binders; absorption without compliance produces reputational risk.

Agentic systems and compliance

Agentic architectures introduce orchestration: models that plan, call tools, retrieve documents, and act across systems. Regulators and customers will ask what the agent was allowed to do and what it actually did.

Design patterns that ease audit:

- **Policy at the orchestrator** (tool allow lists, spend caps, data scopes)
- **Step-level logging** (tool name, arguments hash, outcome, user/session)
- **Human approval before irreversible actions** (payments, external email, record deletion)
- **Evaluation sets** for known failure modes before promotion
- **Rollback** (versioned prompts, corpora, and tool configs)

> **Agentic compliance, put simply:** prove what the system was permitted to do, what it did, and who approved consequential steps, not only that a model answered a chat prompt. **Operating case:** a claims agent may summarise files and draft correspondence automatically, but cannot mark a claim paid without a named adjuster click; every tool call is logged with claim id and timestamp.

Monolithic “send everything to the API” designs can work for low-risk drafting. They become expensive to defend when the same pipeline touches customer eligibility or employment ranking.

The mistake to avoid

Treating the EU AI Act as a legal-only problem that arrives finished from outside counsel.

Engineering hides SaaS features. HR buys hiring AI without informing security. Procurement renews platforms whose AI terms changed silently. The board hears “we are monitoring developments” while twelve undocumented systems run in production.

The CIO who answers all eight questions and assigns owners with dates has done more than most competitors. The general counsel who receives a truthful register can focus on hard classification questions instead of archaeology.

Anti-patterns boards should recognise early

1. **Legal memo without a register.** A thirty-page opinion sits in a folder while nobody maintains a system list. The memo ages; the SaaS stack does not.
2. **“We do not use AI” declaration.** Vendor features count. Shadow consumer tools count when staff use them on company data.
3. **Human oversight as a label.** Slides say “human in the loop” but production runs fully automated ranking with no override log.
4. **Vendor trust without documentation.** “Salesforce/Microsoft handles compliance” is not an artefact. Obtain files; index gaps.
5. **Compliance as a freeze on absorption.** The opposite error also fails: waiting for perfect legal clarity while competitors ship governed workflows. Bounded pilots with gates beat indefinite deferral.
6. **Training as awareness theatre.** Staff need practice on refusal, escalation, and what not to paste into public tools.

Competing interpretation: compliance may lag deployment everywhere

A reasonable counterargument is that enforcement will ramp gradually, national authorities will prioritise large providers first, and mid-market deployers will have time to catch up. That may prove true in ordering. It is a poor bet for customer trust and enterprise sales. EU customers already ask AI diligence questions in security questionnaires. A firm that cannot answer the eight questions will lose deals before it sees a regulator.

Another counterargument is that most tools are minimal risk and the Act is overhead. Many uses are minimal. The failure mode is not over-classifying email summarisation. It is under-classifying hiring, credit, and insurance workflows that were rebranded as “productivity.”

Board decision standard

An AI governance programme clears investment committee when it delivers:

1. **Published register** with at least 90% of known AI-enabled systems listed and owned
2. **Classification memo** for top ten business-impact systems, signed by legal
3. **Human-gate implementation** on every high-impact workflow identified in the memo
4. **Vendor documentation index** with named gaps and chase dates
5. **Incident playbook** with tabletop completed
6. **Training plan** for operators on high-risk and customer-facing paths
7. **Quarterly board reporting** line: new systems, tier changes, incidents, open gaps

This standard is stricter than "we asked legal to watch the news." It is cheaper than an emergency remediation programme after a customer audit fails.

CFO pack: controls that survive scrutiny

Control	Evidence	Cadence
Register completeness	Count of systems vs procurement/IT discovery	Monthly
High-risk workflows gated	List with override logs sampled	Monthly
Vendor doc coverage	% of AI vendors with filed documentation	Quarterly
Incident drills	Tabletop and real incidents reviewed	Quarterly
Training completion	Operators on gated workflows	Per release
Shadow IT findings	Consumer tool use on company data	Quarterly

If the CFO pack cannot be populated, the firm is not ready to claim readiness in customer questionnaires.

Malta and mid-market notes

Malta-origin firms serving EU clients face a dense professional network where reputation travels quickly. A documented AI position is a sales asset in RFPs. An undocumented one is a liability in due diligence.

Small population and talent constraints mean many firms rely on vendor SaaS and external partners. That increases importer/deployer diligence work: you inherit fewer engineering controls and need stronger contract artefacts.

English-language operations and EU membership make Malta a hub for cross-border services (iGaming, fintech, professional services). Multi-jurisdiction customer bases amplify the cost of weak inventory discipline.

FAQ (extended)

Do we need to classify every Copilot seat?

Usually minimal risk for general drafting if data handling is governed and logged. Revisit if Copilot accesses HR files, customer PII, or connects to automated decision workflows.

Is RPA the same as AI under the Act?

Pure rule-based automation without adaptive models may fall outside AI definitions, but many "RPA" products now embed models. Classify by behaviour, not label.

What about open-weight models we host internally?

Hosting can trigger provider or deployer duties depending on modification and placement. Legal analysis required; engineering must still inventory and document. As of July 2026 that inventory must name current weights in use (for example DeepSeek-V4, Zhipu GLM-5.2, Qwen open family), origin of the lab, licence, serve location, and whether any telemetry leaves the EU. Self-hosting Chinese-lab open weights keeps prompts local if you operate the tray; it does not erase AI Act deployer duties or a board-level origin policy.

Does GDPR cover this already?

GDPR and the AI Act overlap on personal data and automated decision-making. Compliance with one does not satisfy the other. Map jointly.

Closing position

The EU AI Act is not the board's reading assignment. It is the board's calendar for accountability while AI deployment accelerates.

Eurostat's 20.0 per cent enterprise adoption figure for 2025 means AI is no longer a edge-case technology in the EU economy. McKinsey's gap between widespread use and rare EBIT impact means boards will scrutinise both value and control. The eight questions in this paper are the minimum operating spine: inventory, tier, data, oversight, transparency, documentation, incidents, vendors.

Answer them with named owners and dates. Legal counsel confirms the hard edges. Engineering and HR make the register true. Procurement makes vendors prove their side. That is readiness in 2026.

References

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act). https://eur-lex.europa.eu/eli/reg/2024/1689/oj
2. European Commission, AI Act implementation and timelines (AI Office). https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
3. Eurostat, "20% of EU enterprises use AI technologies," 11 December 2025 (20.0% in 2025; 13.5% in 2024). https://ec.europa.eu/eurostat/en/web/products-eurostat-news/w/ddn-20251211-2
4. Eurostat, "Use of artificial intelligence in enterprises," Statistics Explained. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises
5. McKinsey & Company / QuantumBlack, "The State of AI: Global Survey 2025." https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
6. Deloitte AI Institute, "The State of AI in the Enterprise: The Untapped Edge" (2026 edition). https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html
7. European Commission, AI Watch (surveys on adoption barriers and enterprise use). https://joint-research-centre.ec.europa.eu/eu-knowledge-services/monitoring-european-digital-policies/ai-watch_en
8. OECD, "Explanatory Memorandum on the Updated OECD Definition of an AI System," 2024. https://www.oecd.org/en/publications/2024/02/explanatory-memo-on-the-updated-oecd-definition-of-an-ai-system_8a9a6fac.html
9. ENISA, "Multilayer Framework for Good Cybersecurity Practices for AI," 2024. https://www.enisa.europa.eu/publications/multilayer-framework-for-good-cybersecurity-practices-for-ai
10. ISO/IEC 42001:2023, Artificial intelligence management system. https://www.iso.org/standard/81230.html

Published by AIMonger . https://aimonger.com . Malta . Europe