Optro

The EU AI Act readiness playbook

Table of contents

Introduction

A global bank based in Amsterdam deploys a new AI e-automation system for retail credit lending decisions. In the rush to deploy the technology across the business and keep up with competitors, the bank fails to implement adequate human oversight and governance controls.

Weeks later, an internal audit reveals a bias in the algorithm: certain demographic groups consistently get denied loans.

Annex III of the EU AI Act classifies systems used to evaluate the creditworthiness of natural persons as high-risk. Deployers of these systems are subject to strict obligations under Article 26, including mandatory human oversight and operational monitoring requirements. Depending on the deployment context and the bank's exact role as deployer or provider, additional obligations — including conformity assessments and a Fundamental Rights Impact Assessment (FRIA) under Article 27 — may also apply.

According to Vision Compliance, 78% of enterprises are unprepared for EU AI Act obligations.

Violations can expose global organizations to severe administrative fines of up to €15 million or 3% of annual global turnover, and, in some cases, even higher penalty tiers.

Despite these significant financial and reputational stakes, many organizations are falling behind. Recent Optro research reveals that only 34% of organizations have an AI model inventory, an explicit EU AI Act obligation for high-risk systems.

This playbook breaks down the Act's core requirements and organizational obligations into a practical reference guide tailored for GRC and audit teams. Use it to assess where your organization currently stands, identify critical control gaps, and build an actionable compliance roadmap.

The EU AI Act compliance is applicable beyond Europe

While the Act is an EU regulation, its reach extends well beyond European borders. Any organization that provides or deploys AI systems used by people in the EU must comply, regardless of where it is headquartered.

EU AI Act enforcement timeline

Understanding the regulatory timeline is the first step in building your compliance strategy. The Act is being enforced in phases, giving global enterprise organizations specific windows to operationalize their controls. Leaders should note that some compliance deadlines will be pushed back as part of negotiations over the Digital Omnibus on AI.

1

February 2025:

Prohibitions on unacceptable-risk AI systems officially go into force. Additionally, requirements regarding AI literacy for staff apply.

2

August 2025:

General-purpose AI (GPAI) model obligations become applicable. (Applies to GPAI model providers).

3

December 2026:

Article 50 transparency obligations take effect for all covered AI systems, including limited-risk systems, per the Omnibus agreement reached on May 7, 2026.

4

December 2027:

Obligations for stand-alone high-risk AI systems defined in Annex III become applicable, per the Omnibus provisional agreement.

5

August 2028:

Obligations for high-risk AI systems used as safety components of products become applicable, per the Omnibus provisional agreement.

How the EU AI Act works

The EU AI Act establishes a detailed legal framework designed to ensure that AI systems deployed within the European Union are safe, ethically sound, and respect fundamental rights. Understanding the foundational elements of this framework is critical to build effective compliance programs.

Scope and territorial reach

The Act casts a wide net, applying to:

Providers: Organizations that develop AI systems or place them on the market

Deployers: Organizations that use AI systems in a professional context

Importers: EU-based entities that bring AI systems from third countries to market

Distributors: Organizations that distribute AI systems in the EU

Crucially, these requirements apply regardless of where an organization is headquartered. If your AI system is used in the EU or affects people there, your organization must comply.

The risk-based classification system

The regulation introduces a risk-based categorization of AI systems. This approach ensures that the most dangerous applications face rigorous scrutiny, while lower-risk applications face minimal regulatory burden. The categories are:

Unacceptable risk: These systems are strictly prohibited.

High risk: These systems face stringent compliance requirements, including risk management, data governance, and conformity assessments.

Limited risk: These systems primarily face transparency obligations.

Minimal or no risk: These systems operate largely without mandatory restrictions, though voluntary codes of conduct are encouraged.

High-risk system categories under Annex III

For most enterprise organizations, high-risk systems represent the heaviest compliance burden. Annex III outlines specific use cases that trigger high-risk obligations. These include biometric identification, the management and operation of critical infrastructure, education and vocational training, and employment and worker management (such as candidate sourcing or screening). It also covers systems that manage access to essential private and public services, including credit scoring and benefit eligibility, as well as law enforcement, migration control, and the administration of justice.

Prohibited practices under Article 5

The unacceptable-risk tier represents the Act's hardest line. Certain applications pose such a severe threat to fundamental rights that they are entirely banned. These prohibited practices include subliminal manipulation, the exploitation of vulnerabilities (such as age or disability), and social scoring by public authorities. The Act also bans real-time remote biometric identification in publicly accessible spaces, emotion recognition in workplaces and educational institutions, and biometric categorization used to infer sensitive personal attributes like race or political opinions.

General-purpose AI models

General-purpose AI (GPAI) models, defined under Article 3(63), sit outside the four-tier risk classification framework and are governed under their own separate obligation tier. Providers of these models face transparency requirements. For models that pose a systemic risk at the Union level, additional safety protocols and adversarial testing requirements become mandatory.

ACTIONS TO TAKE:

Determine your risk level

- [ ] Identify which risk tier applies to each AI system your organization uses or deploys
- [ ] Confirm whether any systems fall under Annex III high-risk categories
- [ ] Determine your organization's role: provider, deployer, importer, or distributor

The governance infrastructure

The EU AI Act does not prescribe a single, rigid governance structure. However, effective compliance requires a unified approach. No single team can operationalize these obligations alone. Your organization must lay the groundwork for governance before attempting to satisfy specific technical requirements.

Establishing cross-functional policies and ownership

Successful AI governance requires collaboration across the entire enterprise. You need to establish an AI steering committee or a similar governance body that includes stakeholders from compliance, risk management, internal audit, legal, IT, and core business units. This cross-functional team must clearly define who owns the risk of specific AI deployments and who is responsible for executing conformity assessments.

In practice, however, GRC functions often hold divergent views on who is actually responsible, making explicit ownership assignment harder than it appears. (Read more on closing the AI oversight gap).

Developing realistic human oversight mechanisms

Human oversight is a non-negotiable requirement for high-risk systems under Article 14. However, GRC teams must work with technical leaders to understand what the Act requires and what is actually operationally realistic. Oversight mechanisms cannot be rubber-stamp approvals. Human reviewers must be fully trained, understand the system's limitations, and possess the authority to override algorithmic decisions or halt the system entirely if severe anomalies occur.

Aligning with existing risk frameworks

Do not build your AI compliance program in a vacuum. The most efficient path to readiness involves aligning AI governance with existing, recognized frameworks such as ISO 42001 or the NIST AI Risk Management Framework (RMF). Integrating these standards helps streamline compliance, reduces duplicative work, and embeds AI risk management directly into your broader enterprise risk architecture. You can read more about unifying your approach in our guide on integrating data and AI governance policies.

ACTIONS TO TAKE:

Build your governance foundation

- [ ] Establish a cross-functional AI governance body with defined ownership
- [ ] Document human oversight protocols for all high-risk systems
- [ ] Map your AI governance program to ISO 42001 or NIST AI RMF

What the EU AI Act requires, and how to know you're ready

The following obligations are sequenced in their logical implementation order: classify first, then assess, document, and monitor.

Core organizational obligations

Identify and classify: Locate all AI systems currently in development or deployment. Classify each by risk tier and document the rationale behind every classification decision.

Build an AI inventory: Maintain a centralized, dynamic inventory covering all in-scope systems.

Conduct fundamental rights impact assessments: Complete FRIAs where required under Article 27, particularly for high-risk deployments affecting essential services or human resources.

Complete conformity assessments: Ensure to fully execute and document required conformity assessments before you place high-risk systems into service.

Implement human oversight: Establish robust human-in-the-loop, human-on-the-loop, or human-in-command mechanisms for all high-risk operations.

Maintain technical documentation: Establish clear protocols for technical documentation and maintain automatically generated event logs to ensure traceability.

Establish post-market monitoring: Put continuous post-market monitoring and incident reporting processes in place. The Act requires you to report serious incidents to national competent authorities quickly.

Address third-party vendor risks: Contractually address all third-party and vendor AI obligations. Importers and distributors share the burden of compliance.

The GRC readiness checklist

Use these questions to pressure-test your current compliance program:

- ☐ Have you identified and classified all internal and third-party AI systems by risk tier?
- ☐ Does a formal, documented rationale for each classification decision exist?
- ☐ Have you completed FRIAs for all high-risk systems?
- ☐ Have you completed conformity assessments prior to the deployment of any high-risk tools?
- ☐ Are human oversight mechanisms clearly defined, operational, and adequately staffed for high-risk systems?
- ☐ Do you actively maintain and securely store technical documentation and operational logs?
- ☐ Is continuous post-market monitoring fully operational?
- ☐ Are third-party AI providers and vendors contractually bound to the Act's compliance obligations?
- ☐ Is there a clearly defined escalation path for reporting and managing severe AI incidents?

Common compliance pitfalls to avoid

Many organizations stumble because they treat AI regulation as a standard IT project rather than a fundamental shift in business operations. GRC teams should watch out for the following traps.

Misclassifying AI system risk levels

Underestimating risk is a costly mistake. Organizations frequently misclassify high-risk use cases — particularly those involving human resources recruitment tools, credit scoring algorithms, and safety-critical operations — as limited risk. A system misclassification like this bypasses the conformity assessments, human oversight requirements, and monitoring controls that high-risk classification would trigger.

As a result, organizations can expose themselves to both regulatory penalties and operational failures they never see coming. That is why rigorous internal review and legal consultation are necessary to validate all classification decisions.

Treating compliance as a one-time project

EU AI Act compliance is not a simple checklist you complete once. AI models drift, underlying data changes, and business use cases evolve. Treating compliance as a one-and-done project rather than an ongoing, continuous monitoring program exposes the organization to regulatory risk.

Building AI governance in a silo

Developing AI policies separate from your existing enterprise risk, data privacy, and IT security frameworks creates operational friction. Governance must be integrated. Privacy, security, and AI ethics often overlap, meaning your governance structure must reflect a unified approach.

The current state of AI governance ownership is worth examining. Optro's report, *The AI oversight gap*, shows that of the GRC and audit professionals recently surveyed, no single function owns more than a quarter of AI governance responsibility: IT (25%), risk management (18%), cross-functional arrangements (17%), and dedicated AI governance teams (10%) all share fragmented control.

Misclassifying AI as a standard IT or vendor risk

Many GRC teams default to managing AI through existing vendor risk or cyber risk workflows, but the EU AI Act imposes obligations that those frameworks were never designed to address. Risk classification, fundamental rights assessments, and conformity assessments have no direct equivalent in a standard vendor management program — treating AI as just another procurement category creates compliance gaps that are difficult to detect and costly to remediate.

Overlooking third-party and vendor obligations

Deployers have explicit obligations regarding transparency, monitoring, and ensuring the tools they purchase meet the Act's standards — the compliance burden does not rest with the provider alone.

Waiting for enforcement to drive action

The timeline for compliance is tight, and building the necessary infrastructure takes months. Organizations that delay action until regulatory authorities begin issuing fines will find themselves scrambling to retrofit complex technical systems, which is both expensive and highly disruptive to business continuity.

ACTIONS TO TAKE:

Audit your current approach

- [ ] Validate all AI risk classifications with legal counsel
- [ ] Audit third-party AI contracts for compliance obligations
- [ ] Confirm management of AI governance as a dedicated program, not absorbed into vendor or cyber risk workflows
- [ ] Schedule regular reviews of AI risk classifications, contracts, and governance program maturity as part of an ongoing effort

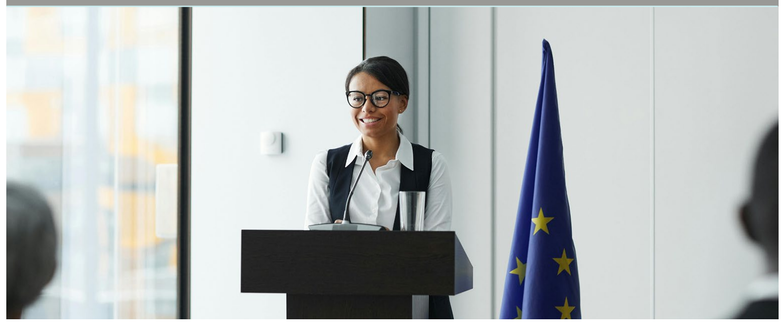

Preparing for the future of AI governance

The Omnibus provisional agreement, reached in May 2026, has extended key enforcement deadlines — providing organizations with additional time to analyze requirements and implement the necessary standards and controls. That window is an opportunity, not a reprieve. Organizations that use it to build cross-functional AI governance infrastructure now will become better positioned as enforcement guidance matures and deadlines arrive.

Keeping pace with the Act's requirements across risk, compliance, internal audit, security, and legal requires a connected view of your AI landscape that is difficult to sustain manually. Explore your AI governance options and see how Optro supports AI Act compliance.

Ready to close your AI compliance gaps?

See how Optro helps GRC teams get there →

About Optro

Optro (formerly AuditBoard) helps enterprises transform risk into opportunity, redefining GRC through an agentic system of action. More than 50% of the Fortune 500 trust Optro to elevate audit, risk, and compliance in addressing a new era of risk. Optro is top-rated by customers on G2 and was named a Leader in the 2025 Gartner® Magic Quadrant™ for Governance, Risk, and Compliance (GRC) Tools, Assurance Leaders. To learn more, visit optro.ai or book a custom demo.