

WHITE PAPER · AI GOVERNANCE

AI Governance 2026 to 2027

A readiness guide for leadership teams: the EU AI Act after the AI Omnibus, the standards that structure governance, and a 90 day plan to put controls in place.

September 2026

At a glance

Who this is for: CEOs, CIOs, CTOs, risk and compliance leaders, and heads of data deploying or buying AI systems.

AI has moved from experiments to systems that touch customers, employees and decisions. Governance is what allows that move to happen safely, and regulation now sets a timetable for it.

The EU AI Act entered into force on August 1, 2024 and applies in phases. Prohibited practices and AI literacy duties have applied since February 2, 2025, and obligations for general purpose AI models since August 2, 2025. The AI Omnibus, in force since July 27, 2026, moved the rules for high risk systems to December 2, 2027 for systems listed in Annex III and to August 2, 2028 for AI embedded in regulated products. The deadline was deferred, not cancelled.

Transparency obligations under Article 50 apply from August 2, 2026, with a grace period to December 2, 2026 for machine readable marking on systems placed on the market before August 2, 2026. Fines reach up to 35 million euros or 7 percent of worldwide annual turnover for prohibited practices.

Organizations do not need to wait for every rule to settle. ISO/IEC 42001 and the NIST AI Risk Management Framework already provide a practical structure. This guide sets out the timeline, the obligations that matter for most businesses, a governance operating model and a 90 day readiness plan.

Key findings

Deferred, not cancelled

High risk obligations now apply from December 2, 2027 (Annex III) and August 2, 2028 (Annex I). Preparation still takes most of that time.

Some rules already apply

Prohibitions, AI literacy and general purpose AI obligations apply today, and Article 50 transparency applies from August 2, 2026.

Reach beyond the EU

The Act can apply to organizations outside the EU when AI systems are placed on the EU market or their output is used in the EU.

Standards give structure

ISO/IEC 42001 is a certifiable AI management system. The NIST AI RMF organizes risk work into Govern, Map, Measure and Manage.

Inventory comes first

No organization can classify, control or document AI it has not found. An AI inventory is the first deliverable of any program.

Contents

- 1 Why AI governance matters now
- 2 The EU AI Act timeline after the AI Omnibus
- 3 Where your AI systems fall
- 4 What non compliance can cost
- 5 The standards that structure governance
- 6 A governance operating model that works
- 7 A 90 day readiness plan
- 8 Readiness checklist
- 9 How Inventh can help
- 10 Sources

01

Why AI governance matters now

Most organizations now use AI in three ways at once: tools employees adopt on their own, AI features inside software they already buy, and systems they build or commission. Each carries different risks, and each needs an owner.

Governance is not a brake on adoption. It is the set of decisions, roles and controls that let an organization say yes to AI with confidence: which uses are acceptable, who approves them, how they are tested, how they are monitored and what happens when something goes wrong.

Regulation adds a timetable. The EU AI Act is the most comprehensive AI law in force, and its reach extends to organizations outside the EU that sell into it or whose AI output is used there.

02

The EU AI Act timeline after the AI Omnibus

The Act applies in phases. The AI Omnibus amendment entered into force on July 27, 2026 and changed several dates. The table shows the position as of September 2026.

Date	What applies
August 1, 2024	The AI Act enters into force.
February 2, 2025	Prohibited AI practices and AI literacy provisions apply.

Date	What applies
August 2, 2025	Obligations for general purpose AI models, governance rules and penalties apply.
July 27, 2026	The AI Omnibus enters into force, changing several dates and simplifying some requirements.
August 2, 2026	Article 50 transparency obligations apply, including disclosure when people interact with AI and labeling of deepfakes.
December 2, 2026	End of the grace period for machine readable marking of AI generated content, for systems placed on the market before August 2, 2026.
December 2, 2027	Rules for high risk AI systems listed in Annex III apply.
August 2, 2028	Rules for high risk AI embedded in products covered by Annex I apply.

The Omnibus also simplified the AI literacy requirement, with the Commission and Member States taking a stronger role, and extended some measures previously reserved for SMEs to small mid cap companies. Confirm the current text with counsel before relying on any simplification.

03

Where your AI systems fall

The Act sorts AI by risk. Most business uses fall in the lower tiers, but a single system in a sensitive area changes the obligations for that system.

Category	Examples	What it means
Prohibited	Social scoring, manipulative techniques that cause harm, certain biometric uses, and non consensual intimate imagery added by the Omnibus	Not allowed. Stop or redesign.
High risk	Annex III areas such as employment and recruitment, access to education, credit scoring and access to essential services	Risk management, data governance, technical documentation, logging, human oversight, accuracy and security requirements.
Transparency	Chatbots, systems that generate or manipulate images, audio or video	Tell people they are interacting with AI and mark or label generated content.
Minimal	Spam filters, recommendations, most internal productivity tools	No specific obligations beyond general law. Good practice still applies.

Roles matter too. A provider develops an AI system or places it on the market. A deployer uses one under its own authority. Most businesses are deployers of AI they buy and providers of AI they build, so many hold both roles for different systems.

What non compliance can cost

Article 99 sets three tiers of administrative fines. For large organizations the higher of the two amounts applies; for SMEs, including start ups, the lower applies.

Infringement	Maximum fine
Prohibited AI practices (Article 5)	35 million euros or 7 percent of worldwide annual turnover
Most other obligations, including provider and deployer duties and Article 50 transparency	15 million euros or 3 percent of worldwide annual turnover
Supplying incorrect, incomplete or misleading information to authorities	7.5 million euros or 1 percent of worldwide annual turnover

The standards that structure governance

ISO/IEC 42001

ISO/IEC 42001 is the international standard for an AI management system. Like ISO/IEC 27001 for information security, it defines policies, roles, risk assessment, controls and continual improvement, and it can be certified by an accredited body. It suits organizations that want a recognized, auditable structure.

NIST AI Risk Management Framework

The NIST AI RMF is a voluntary framework organized into four functions: Govern, Map, Measure and Manage. Its Generative AI Profile, NIST AI 600-1, adds risks and actions specific to generative AI, such as confabulation, data privacy and information integrity. It is widely used as a practical reference, including outside the United States.

The two work together. ISO/IEC 42001 provides the management system; the NIST AI RMF provides a detailed vocabulary for identifying and treating risk inside it.

A governance operating model that works

- 1 Executive owner**
One accountable leader for AI risk, usually the CIO, CTO or chief risk officer, with authority to approve or stop use cases.
- 2 AI governance group**
A small cross functional group from technology, legal, security, data, HR and the business that meets on a fixed cadence and keeps decisions in writing.
- 3 System owners**
A named owner for every AI system, responsible for its purpose, performance, documentation and incidents.
- 4 Policy and standards**
An acceptable use policy for employees, and design standards for systems you build or buy.

The controls every program needs

- AI inventory: every system, its purpose, owner, vendor, data used and risk category.
- Risk classification against the EU AI Act categories and your own risk appetite.
- Data governance: provenance, quality, personal data handling and retention.
- Human oversight: who reviews outputs, when a person must decide, and how to override.
- Testing and evaluation before release: accuracy, bias, robustness and security, including prompt injection for generative AI.
- Monitoring in production: drift, errors, complaints and usage outside the intended purpose.
- Vendor due diligence: how suppliers train, host and secure models, and what they do with your data.
- Incident response: how AI incidents are reported, investigated and fixed.
- Documentation and records that let you show what you did and why.

A 90 day readiness plan

Period	Focus	Deliverables
Days 1 to 30	Discover	Executive owner named, governance group formed, acceptable use policy issued, first AI inventory complete.
Days 31 to 60	Assess	Every system classified, gap assessment against the EU AI Act and ISO/IEC 42001, priority risks agreed.
Days 61 to 90	Control	Controls in place for priority systems, transparency notices live, staff AI literacy training delivered, monitoring started.

After 90 days the program moves to a steady rhythm: quarterly reviews of the inventory, a release gate for new AI systems and an annual review of policy against regulation and standards.

08

Readiness checklist

- ☐ We have an up to date inventory of AI systems, including AI features in purchased software.
- ☐ Every AI system has a named owner.
- ☐ Each system is classified against the EU AI Act risk categories.
- ☐ We know where we act as provider and where we act as deployer.
- ☐ No system falls into a prohibited category.
- ☐ People are told when they interact with an AI system, and generated content is labeled.
- ☐ Staff who use AI have received training appropriate to their role.
- ☐ An acceptable use policy for AI is published and understood.
- ☐ Personal data used by AI systems has a lawful basis and a retention rule.
- ☐ Systems are tested for accuracy, bias and security before release.
- ☐ Production systems are monitored, with thresholds for action.
- ☐ A person can review and override consequential AI decisions.
- ☐ Vendors have answered our AI due diligence questions in writing.
- ☐ AI incidents have a reporting route and an owner.
- ☐ Documentation would let us explain any system to a regulator or customer.

NEXT STEPS

How Inventh can help

Inventh is an AI and custom software development company. We design, build and run AI systems, custom applications and data platforms, with consulting to decide what is worth building first. The services most relevant to this paper:

AI Consulting

AI readiness assessments, AI strategy and roadmaps, governance models and ongoing AI advisory.

<https://inventh.com/services/ai-consulting/>

AI and Machine Learning Development

AI systems built with evaluation, monitoring and documentation from the start.

<https://inventh.com/services/ai-and-machine-learning-services/>

Advanced Analytics and Business Intelligence

Data readiness, governance and the AI ready data platforms that AI depends on.

<https://inventh.com/services/advanced-analytics-and-business-intelligence/>

About Inventh. Inventh is the trading brand of Great Inventh Private Limited, with offices in the Dallas area (United States), Northampton (United Kingdom) and Kochi (India). Every engagement starts with a written proposal that sets out scope, team, timeline and cost before any work begins.

Talk to us: mail@inventh.com · +1 469 607 0077 · inventh.com/contact-inventh

REFERENCES

Sources

- European Commission: AI Omnibus enters into force. <https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force>
- European Commission AI Act Service Desk: Article 99, Penalties. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-99>
- Gibson Dunn: EU AI Act Omnibus agreement, postponed high risk deadlines and other key changes. <https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/>
- ISO: ISO/IEC 42001 Artificial intelligence management system. <https://www.iso.org/standard/42001>
- NIST: AI Risk Management Framework. <https://www.nist.gov/itl/ai-risk-management-framework>

This paper is general information, not legal advice. Regulatory dates and obligations are stated as published by the European Commission and in the Act as of September 2026. Confirm how they apply to your organization with qualified legal counsel. © 2026 Great Inventh Private Limited. Inventh is a trademark of Great Inventh Private Limited.