

**INTERNATIONAL JOURNAL OF LAW
MANAGEMENT & HUMANITIES**
[ISSN 2581-5369]

Volume 9 | Issue 5

2026

© 2026 *International Journal of Law Management & Humanities*

Follow this and additional works at: <https://www.ijlmh.com/>

Under the aegis of VidhiAagaz – Inking Your Brain (<https://www.vidhiaagaz.com/>)

This article is brought to you for free and open access by the International Journal of Law Management & Humanities at VidhiAagaz. It has been accepted for inclusion in the International Journal of Law Management & Humanities after due review.

In case of **any suggestions or complaints**, kindly contact support@vidhiaagaz.com.

To submit your Manuscript for Publication in the **International Journal of Law Management & Humanities**, kindly email your Manuscript to submission@ijlmh.com.

An Examination of Corporate Compliance Gaps in the Implementation of the EU Artificial Intelligence Act

OGHENEHORO EVI ENI*

ABSTRACT

The Regulation on Artificial Intelligence, commonly known as the EU Artificial Intelligence Act, entered into force in August 2024 and has since been applying to companies in stages that stretch from February 2025 through December 2027 and beyond. This staggered timetable, compounded by a Digital Omnibus agreement that has already pushed back the deadline for full compliance with the rules governing stand-alone high-risk systems from August 2026 to December 2027, has left many companies uncertain about precisely what is required of them and by when. This article examines the compliance gaps that have opened between the obligations the Act imposes on providers, deployers, importers and distributors of artificial intelligence systems and what companies operating in or serving the European Union have actually implemented. Drawing on the regulatory text itself, recent readiness research indicating that a large majority of enterprises still lack a basic inventory of the artificial intelligence systems they use, and the professional and academic commentary that has followed the Act's phased rollout, the article traces gaps across five domains, namely the identification and risk classification of artificial intelligence systems, data governance and algorithmic bias, documentation and traceability, human oversight and artificial intelligence literacy, and post-market monitoring. The article concludes with a set of recommendations organized around enterprise-wide risk mapping, contractual allocation of responsibility across the artificial intelligence supply chain, and integrated board-level governance, intended to help companies close the gap between what the Act requires and what most organizations have so far put in place.

Keywords: Corporate Compliance, EU, Artificial Intelligence

I. INTRODUCTION

When the European Union adopted Regulation (EU) 2024/1689, generally known as the Artificial Intelligence Act (the Act), it created the first comprehensive statutory regime anywhere in the world governing the development, sale and use of artificial intelligence systems.¹ The Act does not apply all at once. It entered into force on 1 August 2024 but takes effect in a series of staggered phases, running from 2 February 2025, when the prohibitions on

* Author is a Student at George Mason University, Fairfax, Virginia, United States.

certain artificial intelligence practices began to apply, through 2 December 2027, when providers and deployers of stand-alone high-risk artificial intelligence systems must achieve full compliance with the Act's Chapter III obligations.² This phased structure was itself substantially revised in 2026, when the Council and the Parliament agreed to a Digital Omnibus package, since adopted as Regulation (EU) 2026/1744, which deferred the compliance deadline for the high-risk systems listed in Annex III from 2 August 2026 to 2 December 2027, and deferred the deadline for high-risk artificial intelligence embedded in regulated products, such as medical devices and machinery, from 2 August 2027 to 2 August 2028.³

A regulatory timetable that has already been revised once, only months before an original deadline was due to take effect, does not lend itself to confident corporate planning. Recent industry research illustrates the consequence. A 2026 readiness report by Vision Compliance, drawing on the firm's assessments of its client base across eight industry sectors, found that 78 percent of the organizations assessed had not taken meaningful steps toward compliance, that 83 percent lacked a formal inventory of the artificial intelligence systems they use sufficient to classify those systems by risk level, that 74 percent had no designated internal owner or governance body responsible for artificial intelligence compliance, and that 61 percent had no process in place for generating the technical documentation that high-risk systems require.⁴ These figures form the empirical starting point for this article.

II. THE REGULATORY ARCHITECTURE OF THE EU ARTIFICIAL INTELLIGENCE ACT

Understanding where corporate compliance falls short requires a clear picture of what the Act actually demands, since much of the difficulty companies face stems not from disagreement

¹ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), 2024 O.J. (L 1689), <http://data.europa.eu/eli/reg/2024/1689/oj>.

² See LEGISCOPE, *EU AI Act Effective Dates: Compliance Phases 2024–2027*, <https://www.legiscope.com/blog/eu-ai-act-effective-dates-phases-2024-2027.html> (last visited Oct. 2, 2026) (describing the entry into force date and the successive phases of application through December 2027); see also Regulation (EU) 2024/1689, *supra* note 1, art. 113.

³ See Ahmed Baladi et al., *EU AI Act Omnibus Agreement: Postponed High-Risk Deadlines and Other Key Changes*, Gibson Dunn (May 27, 2026), <https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/> (describing the deferral of the Annex III compliance deadline to 2 December 2027 and the Annex I deadline to 2 August 2028); Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 Amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI), 2026 O.J. (L 1744), <http://data.europa.eu/eli/reg/2026/1744/oj>.

⁴ See VISION COMPLIANCE, *2026 EU AI Act Readiness Report* (2026), as reported in *Vision Compliance Releases 2026 EU AI Act Readiness Report, Finds 78% of Enterprises Unprepared for Obligations*, Nat'l L. Rev. (Apr. 1, 2026), <https://natlawreview.com/press-releases/vision-compliance-releases-2026-eu-ai-act-readiness-report-finds-78>.

with the Act's goals but from the sheer structural complexity of a regulation that classifies obligations by risk tier, by role in the value chain, and by a compliance timetable that now differs across at least five distinct dates.

A. Evolution and objectives of the EU AI Act

The Act was proposed by the European Commission in April 2021 and, following extended trilogue negotiations shaped in part by the rapid public emergence of general-purpose generative systems after 2022, was formally adopted in 2024. Its stated objectives are to ensure that artificial intelligence systems placed on the Union market are safe and respect fundamental rights, while simultaneously fostering investment and innovation by providing legal certainty across the internal market.⁵ The Act is best understood as a product safety regulation that layers organizational and governance duties, drawn substantially from data protection and financial regulatory practice, onto a conformity assessment framework historically used for physical goods.⁶

B. The risk-based regulatory framework

Rather than regulating artificial intelligence as a single category, the Act sorts systems into four tiers: an unacceptable-risk tier that is prohibited outright, a high-risk tier subject to the most extensive obligations, a limited-risk tier subject to targeted transparency duties, and a minimal-risk tier left essentially unregulated. This tiered approach means that the same underlying technology, for instance a machine learning model used to screen job applications, can face dramatically different obligations depending on the specific context and purpose for which a company deploys it, a design choice that places a significant classification burden on companies rather than on the regulator.

C. Prohibited and high-risk AI systems

Article 5 prohibits a defined list of artificial intelligence practices considered to present an unacceptable risk, including systems that deploy subliminal or purposefully manipulative techniques, systems that exploit the vulnerabilities of persons arising from their age, disability or specific social or economic situation, social scoring, whether carried out by public or private actors, and, subject to narrow exceptions, the use of real-time remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement.⁷ High-risk systems,

⁵ Regulation (EU) 2024/1689, *supra* note 1, art. 1.

⁶ See INT'L ASS'N OF PRIVACY PROS., *EU AI Act: Mapping the Interplays with the GDPR*, <https://iapp.org/resources/article/mapping-interplays-gdpr-eu-ai-act> (last visited Oct. 2, 2026) (describing the Act as a product safety regulation that infuses organizational design responsibilities with a concern for individual rights).

⁷ Regulation (EU) 2024/1689, *supra* note 1, art. 5(1).

defined principally by reference to Annex III, include artificial intelligence used in biometrics, critical infrastructure, education and vocational training, employment and worker management, access to essential private and public services, law enforcement, migration, asylum and border control management, and the administration of justice and democratic processes.⁸ It is this high-risk category, comprising many of the artificial intelligence applications most commonly deployed inside large corporations, that carries the extensive obligations examined in Part III.

D. General-purpose AI models and systemic risk

The Act separately regulates general-purpose artificial intelligence models, meaning models trained with a large amount of data at scale that display significant generality and can competently perform a wide range of distinct tasks. It imposes baseline transparency and documentation duties on the providers of such models, with partial exemptions for certain open-source models, and additional obligations, including adversarial testing and systemic-risk mitigation, on the providers of models presumed to carry systemic risk because the cumulative amount of computation used for their training exceeds a defined threshold, currently set at 10 to the power of 25 floating-point operations.⁹ Because many corporate deployers integrate general-purpose models built by a small number of upstream providers rather than developing their own models, Chapter V of the Act has significant downstream consequences, examined further in Part VI.

E. Territorial and extraterritorial scope of the Act

The Act applies not only to providers established within the Union but to providers established outside the Union that place an artificial intelligence system on the Union market or put it into service there, to deployers of artificial intelligence systems established or located within the Union, and, significantly, to providers and deployers established or located in a third country where the output produced by the system is used within the Union.¹⁰ This extraterritorial reach mirrors the approach the General Data Protection Regulation took in 2018, and it means that companies with no physical presence in the Union can nonetheless fall within the Act's scope simply because the outputs of a system they operate are used by people or organizations located there.

⁸ Regulation (EU) 2024/1689, *supra* note 1, art. 6 & annex III.

⁹ Regulation (EU) 2024/1689, *supra* note 1, arts. 3(63), 51(2), 53 & 55.

¹⁰ Regulation (EU) 2024/1689, *supra* note 1, art. 2(1).

III. CORPORATE OBLIGATIONS UNDER THE EU ARTIFICIAL INTELLIGENCE ACT

The Act does not impose a single set of duties on every company that touches an artificial intelligence system. It instead allocates distinct obligations according to the role a given company plays in relation to a particular system, a structure that itself generates significant compliance complexity for companies that occupy more than one role simultaneously.

A. Obligations of providers, deployers, importers and distributors

A provider, meaning the entity that develops a high-risk artificial intelligence system or has one developed and places it on the market under its own name or trademark, bears the most extensive obligations, set out principally in Article 16.¹¹ A deployer, meaning any entity using an artificial intelligence system under its own authority other than in the course of a personal, non-professional activity, bears a narrower but still substantial set of obligations under Article 26, centred on ensuring that a high-risk system is used in accordance with its instructions and that human oversight is genuinely exercised.¹² Importers and distributors bear verification duties intended to ensure that a system already carries the provider's required documentation and conformity marking before it enters the Union market or is made available further along the supply chain.¹³ Because a single corporate group frequently develops some tools internally while deploying others procured from external vendors, many companies find themselves simultaneously a provider for some systems and a deployer for others, each role carrying a different compliance checklist.

B. Risk and quality management systems

Providers of high-risk systems must establish, implement, document and maintain a risk management system operating as a continuous iterative process throughout the system's lifecycle, identifying and evaluating known and reasonably foreseeable risks and adopting mitigation measures accordingly, and must additionally maintain a quality management system covering matters including regulatory compliance strategy, design control, and post-market monitoring procedures.¹⁴ These are not one-time certification exercises but standing organizational processes that must be actively maintained for as long as a system remains in use.

¹¹ Regulation (EU) 2024/1689, *supra* note 1, arts. 3(3), 16.

¹² Regulation (EU) 2024/1689, *supra* note 1, arts. 3(4), 26.

¹³ Regulation (EU) 2024/1689, *supra* note 1, arts. 23–24.

¹⁴ Regulation (EU) 2024/1689, *supra* note 1, arts. 9, 17.

C. Data governance and data quality requirements

Article 10 requires that training, validation and testing data sets used for high-risk systems be subject to appropriate data governance and management practices, including examination for possible biases likely to affect health, safety or fundamental rights, and requires that such data sets be relevant, sufficiently representative and, to the best extent possible, free of errors and complete in view of the intended purpose.¹⁵ This provision sits close to, but is not identical with, the data minimization and accuracy principles found in data protection law, a relationship examined further in Part VIII.

D. Documentation, transparency and human oversight

Providers must draw up technical documentation before a high-risk system is placed on the market and keep it up to date, must design the system to enable the automatic logging of events over its lifetime, must provide deployers with clear instructions for use, and must design the system so that it can be effectively overseen by natural persons, including through interface tools that allow a human overseer to understand the system's capacities and limitations, to remain aware of the tendency towards automatic over-reliance on its output, and to intervene in or interrupt its operation.¹⁶ Certain limited-risk systems, including those generating synthetic audio, image, video or text content, carry separate transparency obligations requiring disclosure that content has been artificially generated or manipulated.¹⁷

E. Post-market monitoring and incident reporting

Providers must establish a post-market monitoring system proportionate to the nature of the artificial intelligence technologies and risks involved, actively collecting and analysing data on the system's performance throughout its lifetime, and must report serious incidents, meaning incidents or malfunctions that lead to death or serious harm to health, serious and irreversible disruption of the management or operation of critical infrastructure, the infringement of obligations under Union law intended to protect fundamental rights, or serious harm to property or the environment, to the relevant market surveillance authorities within defined statutory timeframes.¹⁸ This obligation extends the compliance burden well beyond the point of deployment and requires companies to maintain monitoring capacity for the entire operational life of a system.

¹⁵ Regulation (EU) 2024/1689, *supra* note 1, art. 10.

¹⁶ Regulation (EU) 2024/1689, *supra* note 1, arts. 11–14.

¹⁷ Regulation (EU) 2024/1689, *supra* note 1, art. 50.

¹⁸ Regulation (EU) 2024/1689, *supra* note 1, arts. 3(49), 72–73.

IV. CORPORATE COMPLIANCE GAPS IN THE IMPLEMENTATION OF THE EU AI ACT

Set against the obligations described in Part III, the readiness data introduced in Part I point to a set of recurring gaps that appear across sectors and across companies of different sizes, rather than isolated failures confined to a small number of laggard organizations.¹⁹

A. AI identification and risk classification gaps

The most fundamental gap precedes substantive compliance altogether. A company cannot apply the correct tier of obligation to a system it has not identified as falling within the Act's scope, yet recent research found that 83 percent of assessed organizations lacked a formal inventory of the artificial intelligence systems in use across their operations, a deficiency that makes accurate risk classification under Article 6 and Annex III effectively impossible.²⁰ This gap is frequently compounded by shadow deployment, meaning the informal adoption of artificial intelligence tools, including general-purpose chat interfaces and features embedded within existing software, by individual business units without central review, leaving compliance and legal functions unaware that a system requiring classification even exists.

B. Data governance, data quality and algorithmic bias

Even where a high-risk system has been correctly identified, satisfying the data governance requirements of Article 10 demands documented evidence of representativeness and bias examination that many companies, particularly those relying on data sets assembled before the Act's data governance requirements were defined, are not well positioned to produce retroactively. Academic commentary examining the interaction between the Act and the General Data Protection Regulation has identified, as a general challenge, the need to meet Article 10's data requirements while respecting the data minimization principle, and has stressed that processing special category data specifically to detect and correct bias, something the Act permits (originally under Article 10(5) and, since the Digital Omnibus, under Article 4a) and Article 10 in some contexts appears to require, is permitted only where strictly necessary and

¹⁹ See Vision Compliance, *supra* note 4; see also COMPLIANCE WEEK, *AI & Compliance Survey 2026: Adoption Is High. Governance and Controls Lag*. (Apr. 30, 2026), <https://www.complianceweek.com/resource/ai-compliance-survey-2026-adoption-is-high-governance-and-controls-lag/> (reporting that more than 83 percent of 193 compliance, ethics, risk and audit leaders surveyed use AI tools, but only about 25 percent have implemented a strong governance framework).

²⁰ Vision Compliance, *supra* note 4.

subject to appropriate safeguards.²¹ That leaves compliance functions to judge for themselves how far bias testing can proceed without separately violating data protection obligations.

C. Documentation, transparency and traceability deficiencies

The 2026 readiness research found that 61 percent of organizations had no process in place for generating the technical documentation that Article 11 requires, including the records of data governance decisions, model performance metrics and human oversight procedures that must accompany a high-risk system before it is placed on the market.²² This gap is particularly acute for systems developed some years before the Act's requirements crystallized, where the original design decisions were never contemporaneously documented and must now be reconstructed after the fact, a process that is time-consuming, resource-intensive and, in some cases, simply impossible where the original development team has since left the organization.

D. Human oversight and AI literacy deficits

Article 14's requirement that human overseers be able to understand a system's capacities and limitations presupposes a baseline of artificial intelligence literacy among staff that many organizations have not yet built. The Digital Omnibus has since softened the original artificial intelligence literacy obligation in Article 4 from a duty to take measures to ensure, to their best extent, a sufficient level of literacy among staff to a duty merely to take measures to support the development of that literacy, citing stakeholder experience that stringent literacy obligations were not suitable for all providers and deployers and the additional compliance burden they created, particularly for smaller enterprises.²³ The change reduces the immediate compliance burden, but it arguably also signals that the original standard was proving difficult to meet in practice. Without genuine literacy, human oversight risks becoming a formality in which a nominal reviewer signs off on a system's output without the capacity to question it meaningfully.

²¹ See Hajo Michael Holtz & Jonas Ledendal, *AI Data Governance: Overlaps Between the AI Act and the GDPR*, 18 L. INNOVATION & TECH. 380 (2026), <https://doi.org/10.1080/17579961.2026.2633677> (identifying as a general challenge the need to meet the data requirements of art. 10 of the AI Act while respecting the data minimization principle in art. 5(1)(c) GDPR, and examining the strict necessity and safeguard conditions on processing special categories of personal data for bias detection and correction); *see also* Regulation (EU) 2024/1689, *supra* note 1, art. 10(5) (deleted by Regulation (EU) 2026/1744, *supra* note 3, art. 1(9)(b), and replaced by art. 4a, inserted by art. 1(6) of that Regulation).

²² Vision Compliance, *supra* note 4.

²³ Baladi et al., *supra* note 3; Regulation (EU) 2026/1744, *supra* note 3, recital 8 & art. 1(5) (replacing art. 4 of Regulation (EU) 2024/1689).

E. Monitoring, auditing and reporting gaps

Post-market monitoring under Article 72 and serious incident reporting under Article 73 depend on infrastructure capable of continuously tracking a deployed system's real-world performance against its intended purpose, yet the same governance deficits that leave companies without a system inventory in the first place tend also to leave them without the monitoring infrastructure needed to detect performance drift or emerging harms after deployment. The 74 percent of organizations reported to lack any designated internal owner for artificial intelligence compliance are, almost by definition, poorly positioned to operate the kind of standing monitoring function that Articles 72 and 73 require.²⁴

V. CORPORATE GOVERNANCE AND ACCOUNTABILITY IN AI COMPLIANCE

The compliance gaps described in Part IV are, at root, governance gaps. A company with a clear internal owner, defined escalation paths and board-level visibility is far better positioned to satisfy the Act's substantive requirements than one lacking these structures, regardless of how sophisticated its underlying artificial intelligence systems may be.

A. Board and senior management responsibility

Recent commentary on board-level oversight of artificial intelligence has observed continuing gaps in direct board engagement with artificial intelligence governance.²⁵ Artificial intelligence governance is better treated as a strategic and legal risk warranting direct board attention, comparable to the attention many boards now give to data protection or financial crime compliance, than as a technology or information security matter delegated entirely to management. Because fines under Article 99 can reach 7 percent of a company's total worldwide annual turnover, a ceiling higher even than that for the most serious infringements of the General Data Protection Regulation, the case for direct board engagement rests on ordinary enterprise risk management principles as much as on the Act's specific text.

B. Fragmentation of corporate responsibility

Within many organizations, responsibility for artificial intelligence compliance is currently split across data protection officers, information security teams, procurement functions and individual business units deploying tools for their own purposes, without any single function holding an integrated view of the company's overall artificial intelligence footprint. This

²⁴ Vision Compliance, *supra* note 4.

²⁵ See Evie Wentink, *AI Oversight at the Board Level: A Mid-Year Reckoning*, Ethical Edge (July 17, 2026), <https://ethicaledge.substack.com/p/ai-oversight-at-the-board-level-a> (describing continued gaps in direct board engagement with artificial intelligence governance).

fragmentation is consistent with the finding that a majority of organizations lack a designated internal owner for artificial intelligence compliance as such, since responsibility distributed everywhere in principle tends, in practice, to be owned nowhere in particular.²⁶

C. Legal, compliance and technology functions

Effective compliance requires close coordination between legal and compliance functions, which understand the Act's requirements but often lack visibility into how systems are actually built and trained, and technology functions, which possess that technical visibility but may not appreciate the legal significance of design choices made well before a system is ever formally classified as high-risk. The more effective way to close this gap is to embed compliance requirements directly into existing software development and procurement workflows, rather than to treat legal review as a final gate applied only once a system is nearly ready for deployment.

D. Internal AI governance and accountability structures

Clear accountability for artificial intelligence remains rare: in one 2026 survey of business leaders, only 5 percent said that coordination and accountability were clear across the artificial intelligence lifecycle, and only 17 percent that governance was embedded by design.²⁷ Because a named accountable structure, even before it is fully resourced, gives compliance work an owner, organizations with a defined artificial intelligence governance function, however small, may be expected to make faster progress on policy adoption and system inventory than organizations without one. This reasoning supports treating governance structure itself, rather than any particular technical control, as the foundational compliance gap that companies should close first.

E. Human oversight and automation bias

Even a well-resourced governance structure cannot fully offset automation bias, the well-documented tendency of human reviewers to defer to an automated system's output rather than genuinely interrogate it, particularly under time pressure or when a reviewer lacks the technical background to identify a plausible-sounding but incorrect result. Article 14 anticipates this risk by requiring that human oversight measures specifically address the danger of automatic over-reliance on a system's output, but satisfying this requirement in substance, rather than merely

²⁶ Vision Compliance, *supra* note 4.

²⁷ See ONETRUST, *The OneTrust 2026 AI-Ready Governance Survey Report* (2026), <https://www.onetrust.com/resources/onetrust-2026-ai-ready-governance-report/> (reporting that only 5 percent of surveyed organizations say coordination and accountability are clear across the AI lifecycle, and only 17 percent that governance is embedded by design).

inserting a human into the workflow as a formality, remains one of the harder governance problems the Act poses.²⁸

VI. AI SUPPLY CHAINS, THIRD-PARTY SYSTEMS AND GENERAL-PURPOSE AI

Very few companies build the artificial intelligence systems they use entirely from first principles. Most combine components, models and platforms sourced from multiple external vendors, a reality that the Act addresses directly but that nonetheless generates some of the most persistent compliance gaps this article identifies.

A. Allocation of responsibility across AI supply chains

Article 25 addresses the allocation of responsibility along the artificial intelligence value chain, providing that a distributor, importer, deployer or other third party is considered a provider, and takes on a provider's obligations, where it puts its own name or trademark on a high-risk system already on the market, makes a substantial modification to a high-risk system already on the market in such a way that it remains high-risk, or modifies the intended purpose of a system not originally classified as high-risk in a way that makes it high-risk.²⁹ This provision means that a company that customizes or fine-tunes a third-party system for its own use can unexpectedly find itself reclassified as the system's provider, bearing the full weight of Article 16 obligations it may not have anticipated when it began the customization.

B. Corporate reliance on third-party AI systems

Companies that deploy, rather than build, high-risk systems bear deployer obligations under Article 26 regardless of how little visibility they have into the system's internal design, training data or validation process, a structural asymmetry that leaves many deployers dependent on the provider for information the deployer itself has no independent means of verifying. Where the provider is uncooperative, slow, or itself uncertain about its own compliance position, the deployer's own compliance becomes hostage to a relationship it cannot fully control.

C. Contractual allocation of compliance obligations

Because statutory obligations under the Act cannot be waived or reassigned by private agreement, companies can use contracts with artificial intelligence vendors to allocate the practical burden of compliance, as the Act itself contemplates for high-risk systems,³⁰ for

²⁸ Regulation (EU) 2024/1689, *supra* note 1, art. 14(4)(b).

²⁹ Regulation (EU) 2024/1689, *supra* note 1, art. 25(1).

³⁰ See Regulation (EU) 2024/1689, *supra* note 1, art. 25(4), as amended by Regulation (EU) 2026/1744, *supra* note 3, art. 1(12)(b) (requiring the provider of a high-risk AI system and any third party that supplies an AI system, AI model, tools, services, components or processes used or integrated in it to specify by written agreement the information, capabilities, technical access and other assistance the provider needs to comply).

instance by requiring a vendor to warrant that technical documentation will be kept current, to commit to specified incident notification timeframes, and to provide audit rights sufficient to verify data governance practices. These contractual mechanisms cannot substitute for the underlying statutory allocation of responsibility, but they can meaningfully reduce the information gap that deployers otherwise face.

D. General-purpose AI and downstream corporate users

Companies that build applications on top of a general-purpose artificial intelligence model provided by a third party inherit a layered compliance problem, since the underlying model's provider bears its own Chapter V obligations under Articles 53 and 55, while the downstream company, to the extent its application meets the definition of a high-risk system, separately bears Chapter III obligations of its own and must rely on documentation that the upstream model provider is required, but not always practically incentivized, to supply in a timely and sufficiently detailed manner.³¹ To the extent that a small number of general-purpose model providers supply the underlying technology behind a large share of corporate artificial intelligence applications, this downstream dependency represents a significant systemic compliance vulnerability rather than an isolated contractual issue.

E. Vendor due diligence and continuing compliance monitoring

Closing the supply chain gap requires due diligence that extends beyond the point of initial procurement, since a vendor's compliance posture, and indeed a system's underlying model, can change after a contract is signed through updates, retraining, or a change in the vendor's own regulatory status. Continuing vendor monitoring for artificial intelligence procurement appears to lag behind the processes many companies already operate for cybersecurity or data protection vendor risk, and third-party artificial intelligence risk has itself been described as the single biggest blind spot in enterprise artificial intelligence governance, with only 36 percent of boards said to discuss it regularly.³²

VII. ENFORCEMENT, LIABILITY AND CORPORATE EXPOSURE

The practical significance of the compliance gaps identified above depends heavily on how the Act is enforced, and the enforcement architecture, like the substantive compliance timetable, has itself continued to evolve as the Act's phased application has proceeded.

³¹ Regulation (EU) 2024/1689, *supra* note 1, arts. 25(4), 53(1)(b), 55.

³² See Danny Manimbo, *Enterprise Reality: Why Organizations Aren't as Prepared for AI Governance as They Think They Are*, Cloud Sec. All. (Sept. 16, 2026), <https://cloudsecurityalliance.org/blog/2026/09/16/enterprise-reality-why-organizations-aren-t-as-prepared-for-ai-governance-as-they-think-they-are> (describing third-party AI risk as "the single biggest blind spot in enterprise AI governance today" and reporting that only 36 percent of boards regularly discuss it).

A. Institutional enforcement under the EU AI Act

Enforcement responsibility is divided between a newly created European Artificial Intelligence Office, which supervises general-purpose artificial intelligence models directly at Union level, and national market surveillance authorities designated by each Member State, which supervise high-risk systems within their own jurisdiction, an arrangement that mirrors, without exactly replicating, the national authority structure long familiar from data protection enforcement under the General Data Protection Regulation.³³ The Digital Omnibus further strengthened the Office's supervisory and enforcement powers, suggesting that centralized enforcement capacity, at least with respect to general-purpose models, will continue to grow.³⁴

B. Administrative fines and sanctions

The Act's penalty structure, set out in Article 99, establishes three tiers: a maximum fine of €35 million or 7 percent of total worldwide annual turnover, whichever is higher, for violations of the Article 5 prohibitions; a maximum of €15 million or 3 percent of total worldwide annual turnover for violations of most other substantive obligations, including those governing providers, deployers, importers, distributors and transparency duties; and a maximum of €7.5 million or 1 percent of total worldwide annual turnover for supplying incorrect, incomplete or misleading information to notified bodies or national competent authorities, with lower ceilings applying to small and medium-sized enterprises.³⁵ These figures place the Act's maximum exposure on a scale comparable to, and in the case of prohibited practices exceeding, the penalty regime long familiar from data protection enforcement.

C. Enforcement challenges for multinational corporations

A multinational corporation operating across several Member States faces the practical difficulty of coordinating compliance with national market surveillance authorities that may interpret shared obligations somewhat differently, a difficulty compounded where the same underlying artificial intelligence system is deployed in materially different national regulatory or sectoral contexts. This decentralization risk is not unique to the Act, having been observed previously in cross-border General Data Protection Regulation enforcement, but it is arguably sharper here given the additional technical complexity involved in assessing a system's risk classification and conformity.

³³ Regulation (EU) 2024/1689, *supra* note 1, arts. 64, 70, 74, 88.

³⁴ Baladi et al., *supra* note 3; Regulation (EU) 2026/1744, *supra* note 3.

³⁵ Regulation (EU) 2024/1689, *supra* note 1, art. 99(3)–(6).

D. Evidentiary, technical and audit challenges

Enforcement of a technically complex regime such as the Act requires regulators to develop genuine capacity to audit model behaviour, training data provenance and system logs, a capacity that is still being built within both the European Artificial Intelligence Office and national authorities, which for many high-risk systems the Act assigns to existing product safety market surveillance authorities or to data protection supervisory authorities, bodies whose expertise lies in those adjacent fields rather than in artificial intelligence engineering specifically.³⁶ This capacity gap cuts in a complex direction, since it may reduce near-term enforcement intensity even as it increases long-term uncertainty about how technical concepts such as substantial modification or systemic risk will ultimately be interpreted once regulators develop deeper technical fluency.

E. Corporate liability and regulatory deterrence

Whether the Act's penalty structure will function as an effective deterrent depends significantly on enforcement consistency and visibility, since a maximum fine that is rarely or unpredictably applied provides weaker practical deterrence than a moderate fine applied with consistency and public visibility, a lesson that enforcement experience under the General Data Protection Regulation arguably illustrates. Because the high-risk compliance deadline for Annex III systems has itself only recently been deferred to December 2027, meaningful enforcement data specific to Chapter III obligations remain limited, leaving the deterrent effect of the Act's penalty structure, for now, more a matter of statutory design than of observed regulatory practice.

VIII. REGULATORY OVERLAP AND STRATEGIES FOR CLOSING CORPORATE COMPLIANCE GAPS

The Act does not operate in isolation. It sits alongside an existing and still developing body of European digital regulation, and much of the practical compliance burden companies face is likely to stem from reconciling overlapping, and at times imperfectly aligned, obligations rather than from any single regulation considered on its own.

A. Interaction between the EU AI Act and the GDPR

The Act and the General Data Protection Regulation intersect most directly around data governance, bias testing that requires the processing of special category personal data, and the overlapping requirement for impact assessments, namely the fundamental rights impact

³⁶ See Regulation (EU) 2024/1689, *supra* note 1, art. 74(3), (8).

assessment that the Act requires of certain deployers and the data protection impact assessment that the General Data Protection Regulation already requires for high-risk processing.³⁷ The 2026 readiness research found that organizations already compliant with the General Data Protection Regulation were better positioned for compliance with the Act, particularly in data governance, impact assessments and documentation, though the Act's conformity assessment and post-market monitoring obligations remain additional requirements without a direct analogue in data protection law.³⁸

B. Cybersecurity, digital services and product regulation

Article 15's requirement that high-risk systems achieve an appropriate level of accuracy, robustness and cybersecurity throughout their lifecycle operates alongside separate cybersecurity obligations that apply to many of the same companies under the NIS 2 Directive and, for financial sector entities, the Digital Operational Resilience Act, while artificial intelligence systems that also qualify as online platforms or search engines can additionally trigger obligations under the Digital Services Act.³⁹ For companies operating across these overlapping regimes, much of the additional burden is likely to come from inconsistent terminology and separately maintained compliance documentation rather than from any substantive conflict in underlying policy goals.

C. Enterprise-wide AI risk mapping and compliance by design

The most direct response to the identification and classification gap described in Part IV is a comprehensive, continuously updated inventory of every artificial intelligence system in use across an organization, mapped against the high-risk categories in Annex III and updated whenever a new system is procured, built or materially modified. Embedding this classification step directly into procurement and software development approval workflows, rather than treating it as a separate compliance exercise conducted after a system is already in use, converts compliance from a retrospective audit exercise into a design-stage discipline, an organizational

³⁷ INT'L ASS'N OF PRIVACY PROS., *supra* note 6 (identifying bias monitoring, impact assessments and human oversight as principal points of overlap between the two regimes); *see also* Regulation (EU) 2024/1689, *supra* note 1, art. 27; Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), art. 35, 2016 O.J. (L 119) 1.

³⁸ Vision Compliance, *supra* note 4.

³⁹ Regulation (EU) 2024/1689, *supra* note 1, art. 15; *see also* Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on Measures for a High Common Level of Cybersecurity Across the Union (NIS 2 Directive), 2022 O.J. (L 333) 80; Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on Digital Operational Resilience for the Financial Sector, 2022 O.J. (L 333) 1; Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services (Digital Services Act), 2022 O.J. (L 277) 1.

counterpart to the compliance by design that commentators have read into the Act's requirements for high-risk systems.⁴⁰

D. Third-party due diligence, monitoring and auditing

Closing the supply chain gaps described in Part VI requires standardized vendor due diligence questionnaires specific to artificial intelligence procurement, contractual audit rights sufficient to verify ongoing compliance rather than only compliance at the point of initial contracting, and a continuing monitoring cadence that treats a vendor's artificial intelligence compliance posture as a living risk factor rather than a static attribute confirmed once during onboarding.

E. Integrated corporate AI governance and board-level accountability

Ultimately, the compliance gaps this article has identified converge on a single structural recommendation, namely that companies need an integrated governance function, with clear board-level visibility, that owns the artificial intelligence inventory, coordinates legal, compliance and technology functions, and maintains continuing oversight of both internally built and externally procured systems throughout their operational lifetime, rather than treating artificial intelligence compliance as a project with a fixed completion date tied to the Act's shifting statutory deadlines.

IX. RECOMMENDATIONS

Companies should complete a comprehensive inventory of every artificial intelligence system in current use, whether built internally or procured from a third party, and classify each system against Article 6 and Annex III before undertaking any further compliance work, since the readiness data reviewed in this article suggest that this foundational step remains incomplete for the large majority of organizations.

They should also designate a single accountable governance function for artificial intelligence compliance, with a clear reporting line to senior management and periodic reporting to the board, rather than leaving responsibility distributed informally across legal, compliance, information security and individual business units.

Furthermore, they should embed artificial intelligence risk classification and compliance review directly into procurement and software development approval workflows, so that classification occurs at the point a system is first proposed rather than only once it is already deployed and generating outputs relied upon by the business.

⁴⁰ See Holtz & Ledendal, *supra* note 21 (reading art. 10(2)(a) of the AI Act as requiring compliance by design, an idea on which almost all of the obligations in arts. 8–15 appear to be based and which draws on the principle of data protection by design and by default in art. 25 GDPR).

Finally, they should adopt standardized, artificial intelligence-specific vendor due diligence and contractual audit provisions for every third-party and general-purpose artificial intelligence system they procure, paired with a continuing monitoring cadence rather than a one-time review conducted only at the point of initial contracting.

X. CONCLUSION

The EU Artificial Intelligence Act sets out a detailed and, in principle, coherent framework for ensuring that artificial intelligence systems used within the European Union are safe, transparent and subject to meaningful human oversight. The gap this article has examined is not a gap in the Act's design so much as a gap in corporate readiness to meet a regulatory timetable that has itself proven more fluid than companies may have expected, and in organizational structures that have not yet caught up with the scale of governance the Act presupposes. The evidence reviewed here, from the persistence of incomplete system inventories to the absence of designated governance ownership in a majority of surveyed organizations, suggests that the central challenge facing most companies is not technical sophistication but organizational discipline, namely the unglamorous work of building an inventory, naming an owner, and maintaining documentation continuously rather than reconstructing it under deadline pressure.

The Digital Omnibus deferral of the Annex III compliance deadline to December 2027 offers companies additional time, but that time is valuable only if used deliberately to close the specific gaps this article has identified, rather than treated as a reason to defer engagement further. Given the scale of the penalties the Act authorizes, and the reputational stakes that have already attached to artificial intelligence governance failures more broadly, the companies that use the intervening period to build durable governance structures, rather than undertake last-minute documentation exercises, are likely to be the ones best positioned when full enforcement of the Act's high-risk obligations finally arrives.
