

AIRG INSTITUTE

AI Risk & Governance Institute

Independent Research | Standards Alignment | Governance Frameworks

W H I T E P A P E R . N O . 2

Navigating the EU AI Act:

An Enterprise Compliance Framework for the
World's First Binding AI Regulation

Richmond Perkins Asante

Founder & Executive Director, AIRG Institute

ISO/IEC 42001 Lead Auditor | ISO/IEC 27001 Lead Auditor

Published by AIRG Institute | Columbus, Ohio | 2025

ISSN: Pending | airginstitute.org | info@airginstitute.org

Aligned with: **EU AI Act 2024** | **ISO/IEC 42001:2023** | **ISO/IEC 27001:2022** | **NIST AI RMF**

Executive Summary

The European Union Artificial Intelligence Act (EU AI Act), which entered into force on 1 August 2024 and applies in full from 2 August 2026, represents the world's first comprehensive, binding legal framework for artificial intelligence. It is not a voluntary standard or soft-law guidance ; it is enforceable regulation, with penalties of up to €35 million or 7% of global annual turnover for the most serious violations.

For enterprise organizations deploying AI systems in, or with effects upon, EU markets, the EU AI Act demands immediate and sustained compliance action. Yet many enterprises are ill-prepared: they lack the AI inventory, documentation infrastructure, risk classification methodologies, and governance accountability structures that the Act requires.

This white paper, published by the AI Risk & Governance Institute (AIRG Institute), provides enterprise compliance officers, AI governance leaders, legal and risk functions, and board-level stakeholders with a rigorous, actionable framework for navigating EU AI Act compliance. It maps the Act's requirements onto enterprise governance practice, identifies integration pathways with ISO/IEC 42001:2023 and NIST AI RMF, and provides a phased implementation roadmap calibrated to organizational readiness.

KEY COMPLIANCE ALERT

The EU AI Act's prohibition of certain AI practices (Article 5) applied from 2 February 2025. High-risk AI system requirements apply from 2 August 2026. General-purpose AI model obligations apply from 2 August 2025. Organizations must act now, before enforcement begins.

AIRG Institute's position is unambiguous: EU AI Act compliance is not merely a legal obligation. It represents an opportunity to build the institutional AI governance infrastructure that responsible enterprise AI deployment requires. Organizations that approach compliance strategically, by integrating it with ISO/IEC 42001 certification and established risk management practice, will emerge stronger, more competitive, and better trusted.

1. Understanding the EU AI Act

1.1 Legislative Architecture and Scope

The EU AI Act (Regulation (EU) 2024/1689) establishes a comprehensive legal framework governing the development, deployment, and use of artificial intelligence systems within the European Union. Its jurisdictional reach is explicitly extraterritorial: it applies to any organization that places AI systems on the EU market or whose AI systems produce outputs that affect persons located in the EU, regardless of where the organization is headquartered.

This extraterritorial scope means that enterprises headquartered in the United States, United Kingdom, Asia, or elsewhere are subject to the Act's requirements if their AI systems interact with, affect, or are deployed by EU-based users, employees, customers, or public bodies. For multinational enterprises, the EU AI Act is a global compliance obligation, not a European regional matter.

The Act applies to three categories of actors:

- **Providers:** Organizations that develop AI systems and place them on the EU market or put them into service in the EU, including organizations that develop AI for their own internal use
- **Deployers (Operators):** Organizations that use AI systems under their authority, including enterprises that deploy AI systems procured from third-party providers
- **Importers and Distributors:** Organizations in the supply chain between providers and deployers

Critically, large enterprise organizations frequently occupy both the Provider and Deployer roles simultaneously, developing some AI systems internally while deploying others procured externally. Each role carries distinct compliance obligations that must be mapped and managed separately.

1.2 Application Timeline

The EU AI Act implements a phased application schedule that enterprises must track with precision:

- **2 February 2025:** Prohibition of unacceptable-risk AI practices (Article 5), covering social scoring, real-time biometric identification in public spaces, and subliminal manipulation systems. These prohibitions apply immediately and without transition.
- **2 August 2025:** General-purpose AI (GPAI) model obligations apply (Title VIII). Organizations developing or deploying large language models and other foundation models must comply with transparency, documentation, and safety evaluation requirements.
- **2 August 2026:** High-risk AI system requirements apply in full (Title III). This is the most extensive compliance milestone for most enterprises, covering risk management, data governance, technical documentation, human oversight, and conformity assessment.
- **2 August 2027:** High-risk AI systems that are safety components of products already regulated under EU product safety law gain an extended transition period.

TIMELINE IMPERATIVE

The 2026 application date for high-risk requirements is not a distant horizon ; it is an imminent compliance deadline. Conformity assessments, technical documentation, and registration in the EU AI Act database must be completed before that date. Organizations that begin compliance work in 2026 will be too late.

1.3 Enforcement and Penalties

The EU AI Act establishes a tiered penalty structure designed to create genuine deterrence:

- Up to €35 million or 7% of global annual worldwide turnover (whichever is higher): Violations of the prohibited AI practices in Article 5
- Up to €15 million or 3% of global annual turnover: Violations of high-risk AI system requirements, GPAI obligations, and other Act provisions
- Up to €7.5 million or 1.5% of global annual turnover: Provision of incorrect, incomplete, or misleading information to national authorities

Enforcement is delegated to national market surveillance authorities in each EU Member State, with the European AI Office exercising oversight of GPAI models and coordination of cross-border enforcement. Organizations should not assume that enforcement will be weak or slow. The EU has demonstrated through GDPR that it is willing to impose substantial fines.

2. The EU AI Act Risk Classification Architecture

2.1 The Four-Tier Risk Framework

The EU AI Act's foundational compliance logic is built on a four-tier risk classification system. Every AI system an enterprise deploys must be classified within this framework, because compliance obligations are entirely risk-tier-dependent. Misclassification, such as assigning a high-risk system to a lower tier, constitutes a compliance failure in itself.

| Risk Tier | Definition | Enterprise Examples | Burden |
|--------------|--|--|-----------------|
| Prohibited | Unacceptable risk to fundamental rights | Social scoring by public authorities; real-time biometric surveillance in public spaces; subliminal manipulation | Total ban |
| High Risk | Significant risk to health, safety or fundamental rights; strictly regulated | Credit scoring; CV-screening tools; medical devices; critical infrastructure; law enforcement AI; border control | Full compliance |
| Limited Risk | Transparency obligations apply; lower regulatory burden | Chatbots; AI-generated content; emotion recognition (non-high-risk contexts); deepfake generation | Transparency |
| Minimal Risk | No specific EU AI Act obligations beyond general product safety laws | Spam filters; AI-powered video games; basic recommendation systems; search engine optimization | Minimal |

2.2 High-Risk AI Systems: Annex III Categories

The most consequential classification decisions for most enterprises concern Annex III of the Act, which lists the categories of AI systems designated as high-risk by virtue of their use case and deployment context, regardless of technical design. Enterprises must examine their AI portfolio against each Annex III category with care and legal rigor.

The eight Annex III high-risk categories are:

- **Biometrics:** Remote biometric identification, biometric categorization, emotion recognition systems
- **Critical Infrastructure:** AI used in management of water, gas, heating, electricity, traffic, or digital infrastructure
- **Education and Vocational Training:** AI used to determine access to education, assess students, monitor examinations, or assess learning outcomes
- **Employment, HR and Worker Management:** AI used in recruitment, CV screening, interview evaluation, task allocation, and performance monitoring
- **Access to Essential Private Services and Public Benefits:** AI used in credit scoring, insurance risk assessment, emergency services dispatch, eligibility assessment for public benefits

- Law Enforcement: AI used in crime risk assessment, polygraph systems, evidence reliability evaluation, profiling in criminal investigations
- Migration, Asylum and Border Control: AI used in visa assessment, asylum applications, border control risk assessment
- Administration of Justice and Democratic Processes: AI used in judicial decision support, electoral and political processes

Two high-risk categories warrant particular attention for enterprise compliance functions. First, employment HR AI (that is, AI used in recruitment, screening, and worker management) is ubiquitous in enterprises and frequently underclassified. Second, access to financial services AI, including credit decisioning, insurance underwriting, and fraud detection, touches the core operations of financial services enterprises and carries full high-risk compliance obligations.

2.3 General-Purpose AI Models

The EU AI Act introduces a dedicated compliance regime for general-purpose AI (GPAI) models, which are large-scale foundation models such as large language models trained on broad data that can be adapted to a wide range of downstream tasks. Enterprises that develop or fine-tune GPAI models are subject to Title VIII requirements, which include:

- Technical documentation: Detailed documentation of model architecture, training data, training methodology, capabilities, and limitations
- Copyright compliance: Policies for respecting intellectual property rights in training data
- Transparency to downstream deployers: Information allowing downstream operators to comply with their own Act obligations
- For GPAI models posing systemic risk: Adversarial testing, incident reporting obligations, and cybersecurity measures

Enterprises that deploy third-party GPAI models (without developing them) as Deployers rather than Providers have lighter obligations, but must still ensure their use cases comply with applicable tier requirements and that they have appropriate agreements with GPAI model providers documenting the provider's compliance posture.

3. High-Risk AI System Compliance Requirements

3.1 The Full Compliance Obligation Matrix

For enterprises with high-risk AI systems (and most large enterprises will have at least several), the EU AI Act imposes a comprehensive set of mandatory requirements that must be designed into AI systems before they are placed on the market or put into service. Compliance is not a post-deployment exercise: it must be built in from the design and development stage.

The following table maps the key high-risk requirements to their Article references and enterprise obligations:

| Requirement | Article | Enterprise Obligation |
|-------------------------------------|-----------------------|---|
| Risk Management System | Article 9 | Documented risk identification, analysis, evaluation and mitigation for the full AI lifecycle |
| Data & Data Governance | Article 10 | Training data quality, relevance, completeness; bias examination; data governance practices documented |
| Technical Documentation | Article 11 & Annex IV | Pre-market technical file; architecture, training methodology, performance benchmarks, intended purpose |
| Record-Keeping & Logging | Article 12 | Automatic event logging; audit trail covering operation period; logs retained minimum 6 months |
| Transparency | Article 13 | Clear instructions for use; limitations disclosed; human oversight capabilities described |
| Human Oversight | Article 14 | Operator must be able to understand, monitor, override or stop the AI system |
| Accuracy, Robustness, Cybersecurity | Article 15 | Performance benchmarked; resilience to errors and adversarial inputs; security measures documented |
| Conformity Assessment | Article 43 | Self-assessment or third-party audit depending on system category; CE marking required for EU market |
| Registration | Article 51 | High-risk AI systems registered in EU AI Act database before being placed on market |
| Post-Market Monitoring | Article 61 | Active monitoring of system performance in deployment; incident reporting obligations |

3.2 The Risk Management System: Article 9

Article 9 of the EU AI Act requires providers of high-risk AI systems to establish, implement, document, and maintain a risk management system that covers the entire lifecycle of the AI system. This is not a one-time risk assessment ; it is an ongoing risk management process with specific procedural requirements.

The Article 9 risk management system must:

- Identify and analyze the known and reasonably foreseeable risks associated with the AI system for health, safety, or fundamental rights
- Estimate and evaluate the risks that may emerge when the system is used in accordance with its intended purpose and under conditions of reasonably foreseeable misuse
- Evaluate risks following the analysis of post-market monitoring data
- Adopt suitable risk management measures addressing identified risks

AIRG Institute's view is that organizations with a well-implemented ISO/IEC 42001 AI management system are significantly better positioned to demonstrate Article 9 compliance. ISO/IEC 42001's §6.1 risk assessment and §8.3 risk management requirements are substantially aligned with Article 9's mandate, and an existing ISO/IEC 42001 risk management system can serve as the documented foundation for EU AI Act conformity assessment.

3.3 Data Governance: Article 10

Article 10 imposes specific and demanding data governance requirements on training, validation, and testing datasets for high-risk AI systems. This is one of the most operationally intensive compliance obligations and one that is frequently underestimated in compliance gap assessments.

Article 10 requires that training, validation and testing data must:

- Be subject to appropriate data governance and management practices, including examination for biases
- Be relevant, sufficiently representative, and to the best extent possible free of errors in respect of the intended purpose
- Have appropriate statistical properties including proportional representation across persons or groups
- Be examined for possible biases that are likely to affect health and safety, have a negative impact on fundamental rights, or lead to discrimination

DATA GOVERNANCE IMPERATIVE

Article 10 compliance requires organizations to document not just what data was used, but why it was appropriate, including the characteristics of persons represented in the data and the examination conducted for bias. This requires retroactive documentation for AI systems already in production that will be brought into compliance scope.

3.4 Human Oversight: Article 14

Article 14 requires that high-risk AI systems be designed and developed in such a way as to allow effective human oversight during the period of use. This is not merely a policy statement ; it is a technical and operational design requirement that must be built into the AI system and its deployment context.

Article 14 specifies that human oversight measures must enable persons responsible for oversight to:

- Fully understand the capacities and limitations of the high-risk AI system and be able to duly monitor its operation
- Be aware of the tendency of automatically relying or over-relying on the output produced by the system (automation bias)
- Be able to correctly interpret the AI system's output, taking into account the characteristics of the system and the tools and methods available
- Be able to decide, in any particular situation, not to use the high-risk AI system, or otherwise disregard, override or reverse the output of the system
- Be able to intervene on the operation of the high-risk AI system or interrupt the system through a 'stop' button or similar procedure

The human oversight requirement has profound implications for AI system design. Systems that produce high-stakes outputs without explainability, that operate at speeds that preclude meaningful human review, or whose interfaces do not provide adequate context for operator judgment, will not satisfy Article 14. This may require fundamental redesign of deployed systems.

4. Integration with ISO/IEC 42001 and NIST AI RMF

4.1 The Case for Standards Integration

The EU AI Act does not exist in isolation. Enterprises that approach compliance as a purely legal exercise, mapping obligations to documentation requirements and stopping there, will produce compliance programs that are fragile, expensive to maintain, and unlikely to satisfy the substantive intent of the regulation.

AIRG Institute's approach positions EU AI Act compliance within a broader, standards-based AI governance architecture. Organizations that have implemented, or are in the process of implementing, ISO/IEC 42001:2023 will find that the vast majority of EU AI Act requirements map directly onto ISO/IEC 42001 clauses. This alignment is not coincidental: the Act's technical requirements were developed with reference to international standards, and ISO/IEC 42001 was designed to support regulatory compliance.

4.2 EU AI Act and ISO/IEC 42001 Alignment

The following table maps the EU AI Act's key compliance requirements to corresponding ISO/IEC 42001 clauses and explains the governance alignment:

| EU AI Act Requirement | ISO/IEC 42001 Clause | Alignment Note |
|---|---|---|
| EU AI Act Article 9 (Risk Management) | ISO/IEC 42001 §6.1 (Risk Assessment) | Both require systematic risk identification, analysis and treatment across the AI lifecycle |
| EU AI Act Article 10 (Data Governance) | ISO/IEC 42001 §8.4 (Data for AI Systems) | Aligned requirements for training data quality, governance, and bias examination |
| EU AI Act Article 12 (Record-Keeping) | ISO/IEC 42001 §7.5 (Documented Information) | Parallel documentation and audit trail obligations |
| EU AI Act Article 14 (Human Oversight) | ISO/IEC 42001 §8.5 (System Operation) | Both mandate human oversight capability and defined operator responsibilities |
| EU AI Act Article 61 (Post-Market Monitoring) | ISO/IEC 42001 §9.1 (Performance Evaluation) | Aligned monitoring and evaluation obligations throughout operational life |
| EU AI Act Article 43 (Conformity Assessment) | ISO/IEC 42001 (Certification) | ISO/IEC 42001 certification provides documented evidence supporting conformity assessments |

An organization with a mature ISO/IEC 42001 AI management system has, in effect, built the foundational governance infrastructure for EU AI Act compliance. The gap between ISO/IEC 42001 compliance and EU AI Act compliance is primarily one of specificity. The Act's requirements are more detailed in certain areas (particularly data governance and conformity assessment) but structurally aligned.

4.3 Integration with NIST AI RMF

For multinational enterprises operating in both EU and US regulatory environments, integrating EU AI Act compliance with NIST AI RMF alignment is both possible and strategically sound. The NIST AI RMF's four core functions (GOVERN, MAP, MEASURE, MANAGE) each map coherently onto EU AI Act requirements:

- **GOVERN:** The NIST AI RMF GOVERN function addresses organizational practices, policies, and accountability structures that align directly with the EU AI Act's governance, human oversight, and provider accountability requirements
- **MAP:** The MAP function's AI system context characterization aligns with the EU AI Act's risk classification requirements and intended purpose documentation obligations
- **MEASURE:** The MEASURE function's AI risk analysis and benchmarking activities support EU AI Act Article 9 risk management and Article 15 accuracy and robustness requirements
- **MANAGE:** The MANAGE function's incident response and risk treatment activities align with EU AI Act post-market monitoring and Article 61 serious incident reporting obligations

AIRG Institute recommends that enterprises adopt a unified governance framework that satisfies all three frameworks simultaneously (EU AI Act, ISO/IEC 42001, and NIST AI RMF) rather than maintaining parallel and potentially inconsistent compliance programs. The marginal cost of alignment is low; the benefit of a single, coherent governance architecture is substantial.

5. EU AI Act Compliance: The Enterprise Roadmap

5.1 Phase 1: Assess and Classify (Months 1-3)

The foundation of EU AI Act compliance is a complete, accurate, and well-documented AI system inventory. Many enterprises cannot currently produce such an inventory. AI systems have been deployed across business units, often without enterprise-wide oversight. The first compliance task is to close this gap.

- Conduct a comprehensive AI system inventory across all business units, functions, and geographies
- Classify each system against the EU AI Act's four-tier risk framework, including Annex III category analysis
- Identify prohibited AI practices (Article 5) and take immediate remediation action for any systems in scope
- Identify GPAI model usage, covering both internal development and external provider deployment
- Map the organization's role for each system (Provider, Deployer, or both)
- Conduct jurisdictional analysis: which systems produce outputs affecting EU persons and are therefore in scope
- Document the inventory in a format that will support ongoing regulatory reporting

5.2 Phase 2: Gap Analysis and Remediation Planning (Months 2-4)

With the AI inventory and classification complete, the organization can conduct a systematic gap analysis against the EU AI Act's requirements for each risk tier and role. This gap analysis should be documented at the system level. A single enterprise-wide gap assessment is insufficient.

- For each high-risk AI system: assess compliance status against all Articles 9-15 requirements
- For each GPAI model in scope: assess compliance against Title VIII requirements
- Prioritize remediation by risk tier, proximity to enforcement deadline, and remediation complexity
- Develop system-level remediation plans with owners, timelines, and resource requirements
- Assess third-party AI provider compliance posture and update vendor contracts to include EU AI Act representations and warranties
- Identify systems that cannot be brought into compliance without fundamental redesign; the organization must then determine whether continued deployment is viable

VENDOR MANAGEMENT IMPERATIVE

Enterprises that deploy third-party AI systems as Deployers cannot outsource compliance. They remain responsible for ensuring systems they deploy satisfy EU AI Act requirements. Contracts with AI vendors must be updated to require compliance representations, ongoing

monitoring data, and technical documentation access.

5.3 Phase 3: Build Compliance Infrastructure (Months 3-9)

EU AI Act compliance is not just about individual AI systems. It requires enterprise-wide compliance infrastructure that can sustain compliance across a growing and evolving AI portfolio.

- Establish an EU AI Act compliance program within the AI governance function, with clear ownership and board-level reporting
- Develop technical documentation templates and governance procedures for high-risk AI systems
- Implement system-level risk management documentation processes (Article 9)
- Build or update data governance controls to satisfy Article 10 requirements for training datasets
- Implement automated logging and audit trail capabilities for high-risk systems (Article 12)
- Design and validate human oversight mechanisms for high-risk AI systems (Article 14)
- Prepare conformity assessment documentation and initiate third-party audits where required
- Register qualifying high-risk AI systems in the EU AI Act database (Article 51)

5.4 Phase 4: Sustain and Mature (Ongoing)

EU AI Act compliance is not a one-time project. The Act requires ongoing compliance disciplines that must be embedded into the organization's AI development, procurement, and deployment processes.

- Implement post-market monitoring programs for all high-risk AI systems (Article 61)
- Establish serious incident reporting processes and EU national authority notification procedures
- Integrate EU AI Act requirements into AI procurement processes for new system acquisition
- Conduct annual compliance reviews and update documentation for system changes
- Monitor EU AI Act implementing acts and Commission guidance for evolving requirements
- Pursue ISO/IEC 42001 certification to establish a documented governance baseline that supports future conformity assessments

6. Board-Level Governance and Accountability

6.1 EU AI Act as a Board-Level Risk

The EU AI Act creates material enterprise risk across legal, reputational, operational, and competitive dimensions, that is properly governed at board level. Enforcement penalties of up to 7% of global turnover represent a potentially significant financial exposure for large enterprises. Prohibition of high-risk AI systems not in compliance could disrupt core business processes. Reputational damage from enforcement action or public disclosure of AI Act violations compounds these financial risks.

Boards of directors should treat EU AI Act compliance as a material risk requiring:

- Regular board-level reporting on EU AI Act compliance status and remediation progress
- Clear executive accountability for the compliance program (typically the Chief AI Officer, Chief Risk Officer, or Chief Compliance Officer)
- Sufficient resource allocation for compliance infrastructure, legal analysis, and technical remediation
- Inclusion of EU AI Act risk in enterprise risk registers and disclosed risk factors (where required by securities law)

6.2 The Role of Independent Governance

AIRG Institute's core position is that AI governance credibility requires independence. The organizations that will navigate EU AI Act compliance most effectively and communicate their compliance posture most convincingly to regulators, customers, and investors are those that have built genuinely independent AI governance structures: governance bodies with real authority, audit functions with real independence, and compliance documentation that reflects actual practice rather than aspirational statements.

The EU AI Act's conformity assessment requirements, in particular, reward organizations with documented, auditable governance processes. An ISO/IEC 42001-certified AI management system provides precisely the external validation that conformity assessments and regulatory engagement require.

AIRG Institute, through its Programs & Advisory function and the Enterprise Responsible AI Council (ERAC), supports organizations in building the independent governance infrastructure that EU AI Act compliance and responsible AI deployment demands.

Conclusion

The EU AI Act is transformative regulation. It establishes, for the first time, a legally binding framework that treats AI governance not as a voluntary aspiration but as a mandatory organizational discipline, with real penalties for non-compliance and real obligations to the individuals whose lives AI systems affect.

For enterprise organizations, the Act demands a clear-eyed compliance program: a comprehensive AI inventory, rigorous risk classification, system-level compliance documentation, robust data governance, functional human oversight mechanisms, and sustained post-market monitoring. It demands supplier management practices that extend compliance obligations through the AI supply chain. And it demands board-level accountability for a material category of enterprise risk.

But the EU AI Act also represents an opportunity. Organizations that build genuinely rigorous AI governance infrastructure, with institutional depth rather than a minimum viable compliance posture, will be better positioned to deploy AI at scale, with regulatory confidence, stakeholder trust, and organizational integrity.

AIRG Institute exists to support that work. Through independent research, governance frameworks aligned with international standards, and the convening power of the Enterprise Responsible AI Council, we are committed to helping enterprises navigate the EU AI Act not just as a compliance requirement, but as a step toward the AI governance future that the stakes demand.

Independent. Rigorous. Accountable.

About AIRG Institute

The AI Risk & Governance Institute (AIRG Institute) is an independent non-profit organization dedicated to advancing responsible AI governance through rigorous research, international standards alignment, and multi-stakeholder engagement. AIRG Institute conducts and publishes independent research on enterprise AI risk and governance, develops practical governance frameworks aligned with ISO/IEC 42001 and ISO/IEC 27001, convenes the Enterprise Responsible AI Council (ERAC), and provides advisory services to organizations navigating AI governance challenges.

AIRG Institute operates with full independence from commercial AI developers and maintains strict conflict-of-interest policies to ensure the integrity of its research and governance work.

Programs & Advisory

AIRG Institute's advisory programs support organizations through EU AI Act gap assessments, ISO/IEC 42001 readiness assessments, high-risk AI system compliance documentation, data governance framework development, and board-level AI governance program design. To discuss advisory engagement, contact info@airginstitute.org.

Enterprise Responsible AI Council (ERAC)

ERAC is AIRG Institute's flagship multi-stakeholder governance initiative, convening enterprise AI leaders, risk and compliance professionals, legal experts, civil society representatives, and academic researchers to develop shared AI governance frameworks and advance responsible AI policy. Organizations interested in ERAC membership should contact info@airginstitute.org.

Contact

info@airginstitute.org | airginstitute.org | Columbus, Ohio, USA

About the Author

Richmond Perkins Asante is the Founder and Executive Director of the AI Risk & Governance Institute. He holds qualifications as an ISO/IEC 42001 Lead Auditor and ISO/IEC 27001 Lead Auditor, and has practiced extensively in AI governance, enterprise risk management, and information security. He is the author of AIRG Institute's foundational publications on enterprise AI governance frameworks.

© 2025 AI Risk & Governance Institute. All rights reserved. This publication may be reproduced for non-commercial purposes with appropriate attribution to AIRG Institute. For permissions and inquiries: info@airginstitute.org

This paper provides general governance and regulatory guidance and does not constitute legal advice. Organizations should consult qualified legal counsel regarding their specific EU AI Act compliance obligations. Regulatory requirements referenced reflect the text of Regulation (EU) 2024/1689 as published in the Official Journal of the European Union.