THE EU AI ACT

A COMPLETE GUIDE FOR ORGANIZATIONS

UNDERSTANDING, IMPLEMENTING, AND OPERATING IN COMPLIANCE WITH EUROPE'S LANDMARK AI REGULATION

RISK-BASED RULES
Understand your obligations across the AI risk spectrum

PRACTICAL IMPLEMENTATION
Step-by-step programs, templates, and checklists

REAL-WORLD INSIGHTS
Case studies, examples, and expert analysis

REGULATORY CONTEXT
How the AI Act interacts with GDPR, the DSA, the Data Act, cybersecurity laws, and more

REGULATION (EU) 2024/1689
EXPLAINED.
SIMPLIFIED.
APPLIED.

ADRIAN REINHARDT

The EU AI Act: A Complete Guide for Organizations

Understanding, Implementing, and Operating in Compliance with Europe's Landmark AI Regulation

Steve Publications

This book is available at
https://leanpub.com/theeuaiactacompleteguidefororganizations

This version was published on 2026-08-11

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

Contents

Understanding, Implementing, and Operating in Compliance with Europe's Landmark AI Regulation

About this book: The EU AI Act is the world's first comprehensive legal framework for artificial intelligence, establishing risk-based rules that will reshape how organizations develop, deploy, govern, and audit AI systems across Europe and beyond. This book translates the full text of Regulation (EU) 2024/1689 into clear, actionable guidance for executives, legal counsel, compliance officers, risk managers, product teams, developers, auditors, and consultants. You will find an article-by-article walkthrough, practical implementation programs with templates and checklists, real-world case studies, and authoritative analysis of how the Act interacts with GDPR, the Digital Services Act, the Data Act, cybersecurity legislation, and sector-specific rules. Whether you are building AI from scratch or buying it from a vendor, this guide is designed to help you understand your obligations, design compliant systems, demonstrate conformity, and build governance that lasts.

Introduction: The Age of Regulated AI

Why This Book Exists Now

In June 2024, the European Union crossed a threshold that had been debated for years. On 13 June 2024, the European Parliament and the Council adopted Regulation (EU) 2024/1689, commonly known as the Artificial Intelligence Act or the EU AI Act. The regulation was published in the Official Journal on 12 July 2024 and entered into force on 1 August 2024 [1, 2]. It is the first comprehensive, horizontal legal framework for artificial intelligence enacted by any major jurisdiction.

This book exists because the era of unregulated AI is over for organizations operating in or affecting Europe. The Act does not merely set aspirational principles. It creates specific, enforceable obligations with defined deadlines and penalties that reach up to 35 million euros or 7 percent of global annual turnover [3]. Those numbers alone are enough to demand attention from boardrooms worldwide. But the real impact runs deeper: the AI Act will reshape how organizations design products, manage risk, govern technology, conduct due diligence on vendors, train employees, and think about accountability for algorithmic decisions.

The regulatory timeline is already underway. Prohibitions on certain AI practices became enforceable on 2 February 2025 [4]. Rules governing general-purpose AI models, including large foundation models, began applying on 2 August 2025 [5]. Transparency obligations and most remaining provisions take effect from 2 August 2026, with some high-risk system requirements deferred by the Digital Omnibus amendments adopted in May 2026 to December 2027 or August 2028 [6, 7]. For many organizations, the window for preparation has already closed.

This book is written for practitioners who need more than headlines and executive summaries. It is designed for people who must determine whether their AI systems fall within the Act's scope, classify them correctly, implement required controls, produce documentation, engage with regulators, integrate compliance into existing governance programs, and defend their decisions in

an audit. If you are responsible for ensuring that your organization uses AI lawfully in Europe, this book is for you.

A Global Regulatory Watershed

The EU AI Act matters even if your company has no physical presence in the European Union. Like the General Data Protection Regulation before it, the AI Act has extraterritorial reach. It applies to any provider or deployer outside the EU whose AI systems produce outputs used within the Union [8]. A software company in the United States offering a hiring-screening tool to European employers is covered. A bank in Asia using AI-driven credit scoring for applicants in Germany is covered. A multinational deploying workplace-monitoring AI across its global operations must comply where those operations touch Europe.

The Act's influence extends beyond direct legal reach. It is already shaping regulatory approaches worldwide. Countries and regions looking at their own AI rules are studying the EU model, often adapting elements of its risk-based structure. Industry standards bodies are aligning with its requirements. Supply chains are demanding compliance evidence from vendors. In short, the EU AI Act is becoming a de facto global benchmark for responsible AI governance.

What makes the legislation distinctive is its architecture. Rather than imposing uniform rules across all artificial intelligence, it uses a tiered approach based on risk:

- **Unacceptable risk:** Certain practices are banned outright because they are deemed fundamentally incompatible with EU values and fundamental rights [9].
- **High risk:** AI systems that can significantly affect health, safety, or fundamental rights face strict obligations covering data quality, transparency, human oversight, accuracy, cybersecurity, conformity assessment, and ongoing monitoring [10].
- **Limited (transparency) risk:** Some AI systems must simply disclose their nature to users so people know they are interacting with a machine rather than a human [11].
- **Minimal risk:** The vast majority of current AI applications face no specific obligations under the Act [12].

This risk-based structure mirrors existing EU product-safety regulation. The same logic governs toys, machinery, medical devices, and many other categories: the greater the potential harm, the stricter the controls. The innovation is applying that proven framework to a technology class that evolves faster than most regulated products have in history.

What You Will Learn in These Pages

This book is organized to take you from foundational understanding through practical implementation and into long-term operational compliance. It is structured so you can read it cover to cover or use it as a reference for specific questions.

The first part establishes context and explains the law itself. You will learn how the AI Act was developed, what motivated its provisions, and why it takes the form it does. You will find precise definitions of key terms from Article 3 that determine whether you are a provider or deployer, importer or distributor, and whether your system qualifies as high risk or general-purpose AI. You will walk through the prohibited practices in Article 5, understand how Annex I and Annex III classify high-risk systems, and see what obligations attach at each tier.

The second part dives into the technical and governance requirements. For high-risk AI systems, you will find detailed explanations of data-governance obligations under Article 10, documentation requirements under Article 11, transparency duties under Article 13, human-oversight design under Article 14, accuracy and robustness standards under Article 15, cybersecurity expectations under Article 16, quality-management systems under Article 17, record-keeping under Articles 18 and 19, and post-market monitoring under Article 72. For general-purpose AI models, you will learn about the obligations in Chapter V, including technical documentation, training-data summaries, copyright compliance, and systemic-risk requirements for the most advanced models.

The third part focuses on implementation. You will find a phased program for building an EU AI Act compliance capability: discovering and inventorying your AI systems, classifying them under the risk framework, assessing gaps against legal requirements, remediating deficiencies, operationalizing controls, validating compliance, deploying or continuing to deploy with confidence, and establishing continuous monitoring and improvement. Each phase

includes step-by-step procedures, decision frameworks, responsibility matrices, document templates, checklists, and practical examples for organizations of different sizes and maturity levels.

The fourth part addresses governance integration and complex scenarios. You will learn how to embed AI Act compliance into existing enterprise risk management, privacy, cybersecurity, model-risk management, quality-management, procurement, internal audit, and corporate-governance programs. Dedicated guidance covers generative AI, third-party and open-source models, foundation models, AI-enabled products, internally developed systems, purchased SaaS tools, employee use of AI, legacy systems, and multinational organizations with complex supply chains. Case studies illustrate how the rules apply in realistic situations, including ambiguous cases where interpretation is still evolving.

Throughout the book, I distinguish clearly between:

- **Statutory requirements** directly imposed by the text of Regulation (EU) 2024/1689
- **Delegated and implementing acts** adopted or expected to be adopted by the European Commission under powers granted by the Act
- **Harmonized standards** developed by CEN-CENELEC that will confer a presumption of conformity once cited in the Official Journal
- **Codes of practice** published by industry bodies or the AI Office that provide practical guidance but do not create independent legal obligations
- **Official guidance** from the European Commission, the AI Office, and national authorities that clarifies interpretation without altering the law
- **Areas where interpretation remains unsettled**, flagged explicitly so you can manage uncertainty rather than assume certainty

Every substantive factual claim in this book is cited to a source. Where sources disagree or guidance is pending, I say so directly. The goal is not to give you comfortable answers to every question but to give you reliable information and a clear sense of where the law is firm versus where it is still taking shape.

The regulatory landscape will continue to evolve after this book goes to print. New delegated acts will be adopted. Standards will mature. National authorities will issue guidance. Courts will interpret provisions. The most important thing you can do is build a governance program that is rigorous enough to satisfy the Act as written, adaptable enough to respond to new

requirements, and documented well enough to demonstrate compliance when regulators ask for evidence.

The chapters ahead show you how to do exactly that.

Chapter 1: Genesis of the EU AI Act

The Political Journey from Proposal to Law

The story of the EU AI Act is not simply a legal history. It is a record of how one of the world's most powerful economies grappled with a technology moving faster than its institutions, and ultimately chose to regulate rather than wait.

The legislative process began on 21 April 2021, when the European Commission published its proposal for a Regulation laying down harmonised rules on artificial intelligence [13]. The proposal arrived at a moment of rising public concern about algorithmic bias, surveillance technologies, and the opaque decision-making of automated systems. It also came as China was advancing its AI strategy, the United States was debating how to govern emerging technologies, and industry stakeholders were pressing for clarity about what rules would apply.

The Commission's initial proposal introduced a risk-based framework with four tiers: unacceptable risk (prohibited), high risk (strict regulation), limited risk (transparency obligations), and minimal risk (no specific rules). This structure was largely preserved in the final text, though the details evolved considerably over three years of negotiations among the Commission, the European Parliament, and the Council of the European Union.

The European Parliament adopted its position on 14 June 2023 [14]. The Council reached a general approach in December 2023 [15]. Intensive trilogue negotiations followed, with a provisional political agreement reached on 9 December 2023 [16]. Formal adoption by the Parliament occurred on 13 March 2024 and by the Council on 9 June 2024. The regulation was signed on 13 June 2024, published in the Official Journal on 12 July 2024, and entered into force on 1 August 2024 [1].

The total text comprises 180 recitals, 113 articles organized into 17 chapters, and 13 annexes. It is a dense piece of legislation that draws on decades of EU regulatory experience in product safety, data protection, consumer rights, and fundamental-rights protection. Understanding its origins helps explain why it takes the form it does and what trade-offs were made along the way.

Competing Visions and Key Negotiations

Three broad visions competed during the AI Act's development: a maximalist approach that would have regulated nearly all AI applications with strict requirements, a minimalist approach that would have relied on existing laws and voluntary standards, and the risk-based middle path that ultimately prevailed.

The Parliament's initial position leaned toward the stricter end of the spectrum. It proposed broader definitions of high-risk systems, more extensive transparency obligations, stronger fundamental-rights safeguards, and higher penalties. Several MEPs advocated for prohibitions on additional practices beyond those in the Commission's original proposal, including certain forms of biometric surveillance and workplace monitoring.

The Council, reflecting the positions of member-state governments with varying economic interests and regulatory philosophies, pushed back on some of these expansions. Several member states were concerned about overregulation that could disadvantage European AI companies competing globally. Others emphasized the need to avoid duplicating existing sector-specific rules in areas like medical devices, automotive safety, and financial services.

Industry stakeholders presented a mixed picture. Large technology companies with established compliance functions generally supported clear rules that would level the playing field. Startups and small enterprises worried about compliance costs that could stifle innovation. Some industry groups advocated for more use of regulatory sandboxes, voluntary codes of practice, and international harmonization to reduce fragmentation.

Civil-society organizations, academic experts, and fundamental-rights advocates pressed for stronger protections. They argued that the risks from AI systems in areas like law enforcement, hiring, credit scoring, and public-benefit allocation were too great to leave to voluntary measures or fragmented national rules. Many called for explicit bans on practices they viewed as inherently harmful, including social scoring, certain forms of predictive policing, and real-time biometric surveillance in public spaces.

The final text reflects a careful balancing of these positions. Several key negotiations shaped the outcome:

General-purpose AI models. The original Commission proposal did not include specific rules for foundation models or large language models. As

the technology advanced during negotiations, Parliament and several member states pushed for a dedicated regime. The compromise created Chapter V, which imposes obligations on providers of general-purpose AI models, including those with systemic risk [17]. This was one of the most significant additions during trilogue negotiations.

Open-source exemptions. Open-source developers and researchers lobbied intensely for an exemption to avoid chilling innovation. The final text includes a limited carve-out: the Act does not apply to AI systems released under free open-source licenses, unless they are prohibited practices, high-risk systems, or subject to transparency obligations [18]. However, the exemption does not extend to general-purpose AI models with systemic risk, regardless of their licensing terms.

Prohibited practices. Parliament sought to expand the list of banned AI practices beyond the Commission's initial proposal. The final Article 5 includes prohibitions on subliminal manipulation, exploitation of vulnerabilities, social scoring, predictive policing based solely on profiling, untargeted facial-image scraping from the internet or CCTV, emotion recognition in workplaces and schools, and certain biometric categorization practices [19]. Real-time remote biometric identification in public spaces by law enforcement is prohibited except under narrowly defined circumstances with judicial or administrative authorization.

Penalties. The Parliament's position increased penalty levels above those in the Commission's original proposal. The final Article 99 sets maximum fines of up to 35 million euros or 7 percent of global turnover for violations of prohibited practices, up to 15 million euros or 3 percent for other infringements, and up to 7.5 million euros or 1 percent for supplying incorrect information [20]. These exceed the GDPR's maximum of 4 percent of turnover, signaling the seriousness with which the EU treats AI-related harms.

Timeline. The Act was originally designed with a phased implementation schedule starting from early 2025 and culminating in full application by August 2027. However, practical challenges in operationalizing certain provisions led to the Digital Omnibus package agreed on 6 May 2026, which deferred some high-risk system deadlines to December 2027 or August 2028 while adding a new prohibition on AI-generated non-consensual intimate imagery [6].

Why Risk-Based Regulation Won Out

The risk-based approach that defines the EU AI Act was not inevitable. Alternatives were seriously considered during the legislative process, and each had supporters among policymakers, industry, and civil society.

A **technology-neutral general law** would have applied uniform rules to all decision-making systems regardless of whether they used AI or traditional algorithms. Supporters argued this would avoid regulatory arbitrage and ensure that similar risks were treated similarly. Critics countered that AI systems pose distinctive challenges: opacity in their internal reasoning, emergent behaviors not anticipated by designers, rapid iteration cycles, and the potential for harm at scale. The Parliament's position reflected growing consensus that a dedicated framework was necessary.

A **sector-by-sector approach** would have left AI regulation entirely to existing laws governing specific industries: medical-device rules for health AI, financial-services regulations for credit-scoring models, employment law for hiring tools, and so on. This approach had the advantage of building on established expertise and avoiding overlap. But it also risked gaps where no sector-specific rule applied, inconsistent standards across sectors, and slow adaptation to new technologies. The Commission's impact assessment found that a horizontal framework would be more effective at addressing cross-cutting risks [13].

A **precautionary ban** on certain categories of AI technology was proposed by some MEPs and civil-society groups, particularly for applications in law enforcement, biometric surveillance, and automated decision-making affecting fundamental rights. While the final Act does prohibit specific practices, it stopped short of broader technology bans. Legislators ultimately concluded that a risk-based approach could capture the worst abuses while leaving room for beneficial uses under appropriate safeguards.

The **risk-based framework** won because it offered something each camp could accept:

- Civil society got prohibitions on clearly harmful practices and strict rules for high-stakes applications.
- Industry got clarity about what was allowed, proportionality in requirements scaled to actual risk, and a single EU-wide standard replacing potential national fragmentation.

- Policymakers got a flexible structure that could adapt as technology evolved, with delegated acts allowing technical updates without full legislative revisions.

The framework also aligned with the EU's existing regulatory DNA. European product-safety law has long used risk classification to determine the stringency of requirements. The GDPR uses similar logic: processing special categories of personal data triggers enhanced protections because the risk is higher. The AI Act extends this playbook to algorithmic systems.

The Act's Core Objectives

The stated objectives of the EU AI Act appear in its recitals and opening articles, and they help explain why specific provisions take their particular form. Four themes dominate:

Protecting fundamental rights. Recital 10 states that the regulation aims to ensure a high level of protection of health, safety, and fundamental rights recognized under Union law [21]. This commitment runs through every tier of the framework. Prohibited practices are those deemed fundamentally incompatible with EU values regardless of context. High-risk requirements are designed to prevent discrimination, ensure accountability for automated decisions, and preserve human agency in critical domains. The Act explicitly references the Charter of Fundamental Rights of the European Union multiple times.

Ensuring legal certainty. Recital 11 emphasizes the need for a single, coherent framework applicable across all member states [22]. Before the AI Act, organizations faced potential fragmentation: different national rules on algorithmic transparency, varying interpretations of existing consumer-protection and anti-discrimination laws, and uncertainty about whether new technologies fell within existing regimes. The Act's harmonization objective means that compliance with EU-wide rules should suffice across all 27 member states, reducing the burden on businesses operating cross-border.

Fostering innovation. Recital 13 acknowledges that regulation should not stifle technological progress but instead create conditions for trustworthy AI to flourish [23]. The Act includes several mechanisms designed to support innovation: regulatory sandboxes where new systems can be tested under

supervision, codes of practice providing practical guidance, exemptions for research and development activities, simplified procedures for small enterprises, and a presumption of conformity for systems meeting harmonized standards. The European Commission has also established the AI Office partly to provide industry with clear signals about expectations [24].

Maintaining competitiveness. Recital 14 notes that a robust regulatory framework can enhance the EU's global standing by demonstrating leadership in responsible AI governance [25]. This is not merely rhetorical. By setting high standards, the EU aims to influence global norms and potentially give European companies a competitive advantage as markets worldwide demand trustworthy AI systems. The Act's extraterritorial reach ensures that non-EU companies must also comply when targeting European users, reducing the risk of a race to the bottom.

These objectives are not always perfectly aligned. Protecting fundamental rights can conflict with fostering innovation if compliance costs are too high. Ensuring legal certainty can clash with maintaining competitiveness if rules become outdated as technology evolves. The Act's design reflects ongoing efforts to balance these tensions through mechanisms like delegated acts, standards development, and regulatory sandboxes.

Understanding these objectives matters for practitioners because they inform how regulators will interpret ambiguous provisions. When faced with unclear language, authorities are likely to consider which interpretation best advances the Act's stated goals. If you are designing an AI system or building a compliance program, keeping these objectives in mind can help you anticipate regulatory expectations and build systems that satisfy both the letter and the spirit of the law.

The next chapter establishes the foundational definitions and scope rules that determine whether your organization and its AI systems fall within the Act's reach at all.

Chapter 2: Scope, Definitions, and Territorial Reach

What Is an AI System Under the Act

Before you can determine what obligations apply to you, you must know whether your technology qualifies as an "AI system" under the law. The definition in Article 3(1) of Regulation (EU) 2024/1689 is deliberately broad, but it has important boundaries [26].

An AI system is defined as a machine-based system that:

- Is designed to operate with varying levels of autonomy
- May exhibit adaptivity after deployment
- For explicit or implicit objectives, infers how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments [27]

This definition captures a wide range of technologies beyond what some people think of when they hear "artificial intelligence." It includes:

- Machine-learning models trained on data and deployed for prediction or classification
- Large language models and other generative AI systems
- Expert systems and symbolic reasoning engines
- Systems that combine traditional software with learning components
- Recommender systems that adapt based on user behavior
- Computer-vision systems for object detection, facial recognition, or scene understanding

The key elements are inference (drawing conclusions from inputs rather than simply executing fixed rules) and the capacity to influence environments (whether by providing information that affects human decisions or by directly

controlling physical or digital processes). A system does not need to be fully autonomous to qualify. Even systems where humans make the final decision can be AI systems if the AI component generates outputs that influence those decisions.

Conversely, the definition excludes some technologies that might intuitively seem like AI:

- Rule-based systems with no learning or inference capability (if they simply follow predetermined if-then logic without adaptation)
- Randomization techniques that do not involve inference
- Purely statistical analyses performed by humans using standard tools
- Hardware components considered as separate products, unless they incorporate an AI system

The definition also distinguishes between an **AI model** and an **AI system**. An AI model is defined in Article 3(2) as the result of training through which an AI system derives parameters enabling it to generate outputs for given inputs [28]. A model is a component; a system is what gets deployed. This distinction matters because different obligations apply to models (particularly general-purpose AI models under Chapter V) versus systems deployed for specific purposes (under Chapters III and IV).

Key Defined Terms Every Practitioner Must Know

Article 3 contains over sixty definitions that structure the entire regulation. Here are the ones you need to understand first, grouped by category.

Market actors:

- **Provider** (Article 3(3)): A natural or legal person, public authority, agency, or other body that develops an AI system or has an AI system developed and places it on the market or puts it into service under its own name or trademark. This includes entities that substantially modify an existing AI system [29].
- **Deployer** (Article 3(4)): A natural or legal person, public authority, agency, or other body using an AI system under its authority, except where the system is used in the course of a personal non-professional activity [30].

- **Importer** (Article 3(6)): An entity established in the Union that places an AI system from a third country on the Union market.
- **Distributor** (Article 3(7)): An entity in the supply chain, other than the provider or importer, that makes an AI system available on the Union market.
- **Authorized representative** (Article 3(5)): A natural or legal person established in the Union explicitly mandated by a provider to act on its behalf regarding specific tasks under the regulation.
- **Operator** (Article 3(8)): A collective term encompassing providers, deployers, authorized representatives, importers, distributors, and product manufacturers [31].

Your role determines your obligations. An organization can be both a provider and a deployer: a bank that develops its own credit-scoring AI is a provider of that system but a deployer of other AI tools it purchases from vendors. The Act assigns different duties to each role.

Market actions:

- **Placing on the market** (Article 3(9)): The first making available of an AI system on the Union market.
- **Making available on the market** (Article 3(10)): Any supply of an AI system for distribution, consumption, or use on the Union market in the course of a commercial activity, whether in return for payment or free of charge.
- **Putting into service** (Article 3(12)): The stage at which an AI system is placed at the disposal of a deployer for operation under conditions determined by the provider.
- **Recall** (Article 3(13)): Any measure aimed at achieving the return of an AI system that has already been made available to the deployer.
- **Withdrawal** (Article 3(14)): Any measure aimed at preventing an AI system in the supply chain from being made available on the market.

These temporal markers matter because obligations trigger at specific points in the lifecycle. A provider's conformity-assessment duties apply before placing a system on the market or putting it into service. Post-market monitoring begins after deployment.

Technical terms:

- **Training data, validation data, testing data** (Article 3(15)-(17)): Data used at different stages of model development to teach the system, tune its parameters, and evaluate its performance respectively. Article 10 imposes specific quality requirements on each [32].
- **Deep fake** (Article 3(18)): Content that falsely appears to be authentic or truthful and is likely to mislead, generated or significantly altered by an AI system [33].
- **Conformity assessment** (Article 3(20)): The process demonstrating whether requirements of the regulation have been fulfilled.
- **CE marking** (Article 3(21)): The mark indicating that a high-risk AI system conforms with applicable Union harmonization legislation, including the AI Act [34].

Governance terms:

- **AI Office** (Article 3(40)): The European Artificial Intelligence Office within the Commission responsible for oversight of general-purpose AI models and coordination of consistent application across member states [35].
- **National competent authority** (Article 3(42)): Authorities designated by each member state to supervise implementation and enforcement.
- **Market surveillance authority** (Article 3(43)): Authorities responsible for checking that AI systems on the market comply with applicable requirements.
- **European Artificial Intelligence Board** (Article 3(41)): A body providing cooperation, coordination, and expertise to ensure consistent application of the Act across the Union [36].

Risk terms:

- **Serious incident** (Article 3(49)): An event resulting in death or serious damage to health, or a disruption of operation of critical infrastructure that has an immediate impact on public health, safety, security, or the economy [37].
- **Systemic risk** (Article 3(50)): Risks arising from high-impact capabilities of general-purpose AI models that could have significant negative effects across the Union market, including on democratic processes, public security, and fundamental rights [38].

These definitions are not merely vocabulary. They are legal triggers. Misclassifying your role, misunderstanding when a system is "placed on the market," or underestimating what constitutes a "serious incident" can mean missing obligations you thought did not apply to you.

Territorial and Material Scope

Article 2 of the regulation establishes where and to whom the AI Act applies [39]. Understanding scope is critical because misreading it in either direction creates risk: applying requirements that do not exist wastes resources, while failing to apply requirements that do exist exposes you to enforcement action.

Territorial scope. The AI Act applies to:

- Providers that place AI systems on the Union market or put them into service in the Union, regardless of where they are established [40].
- Deployers established in the Union, for all AI systems they use within the Union.
- Providers and deployers established outside the Union whose AI systems produce outputs used within the Union [41].

This extraterritorial reach mirrors the GDPR's approach. A company headquartered in California that sells a recruitment-screening AI tool to employers in France must comply with the Act as a provider. A hospital in Germany using an AI diagnostic system developed in Singapore must comply with deployer obligations. The test is not where the technology was built but whether it is made available in or used within the European Union in a professional context.

The regulation also applies to EU institutions, bodies, offices, and agencies when they develop or deploy AI systems [42]. This means public-sector organizations across Europe are subject to the same framework as private companies, though some provisions have specific adaptations for law enforcement and defense contexts.

Material scope. The Act covers all AI systems as defined in Article 3(1), with certain exclusions listed in Article 2. If your system falls within the definition and none of the exclusions apply to you, the Act is relevant. Your next step is then to determine which tier of obligations applies based on the risk classification.

Exclusions and Carve-Outs

Article 2 lists activities and systems that fall outside the regulation's scope [43]. These carve-outs are important because they define the boundaries of what the EU chose not to regulate through this framework.

Military and defense. AI systems exclusively for military or national-defense purposes are excluded [44]. However, dual-use systems with both civilian and military applications remain covered when used in civilian contexts. This exclusion reflects the sensitivity of defense matters and the fact that many member states prefer to handle military technology regulation through separate channels.

Research and development. AI systems developed solely for scientific research and development are excluded [45]. The exemption is narrow: it applies only while the system is genuinely in R&D, not once it is deployed for operational purposes. A university laboratory experimenting with new algorithms benefits from this carve-out; a company using an experimental system to make real decisions about customers does not.

Open-source systems. AI systems released under free and open-source licenses that allow access, use, modification, and redistribution are generally excluded [46]. However, the exemption has important limitations:

- It does not apply if the open-source system falls within prohibited practices under Article 5.
- It does not apply if it is classified as a high-risk AI system under Articles 6 to 7.
- It does not apply to transparency obligations under Article 50 for systems designed to interact with humans or generate synthetic content.
- For general-purpose AI models, the exemption applies only if they do not pose systemic risk [47].

The Act does not define "open-source," leaving room for interpretation. Recital 102 describes it as software and data released under licenses that allow open sharing where users can freely access, use, modify, and redistribute them [48]. Licenses with significant restrictions (such as research-only or non-commercial-use clauses) may not qualify.

Personal non-professional activity. The Act does not regulate AI systems used in the course of purely personal, non-professional activities [49]. A person

using a consumer chatbot for entertainment is not covered. However, any professional use triggers applicability: an employee using generative AI to draft client communications is acting professionally even if the tool itself was designed for consumers.

Specific technical exclusions. Article 2 also excludes certain narrow categories:

- Software functions listed in Annex II that are common software functions (such as spam filters, video games, or basic data-processing utilities) unless they incorporate an AI system subject to separate obligations [50].
- Systems used solely for detecting cybersecurity threats and ensuring their prevention, detection, and analysis.

Law enforcement cooperation with third countries. Public authorities in third countries using AI systems within the framework of international law-enforcement cooperation are excluded if adequate safeguards exist for fundamental-rights protection [51]. This recognizes the complexity of cross-border policing while maintaining standards where possible.

Practitioners should approach exclusions cautiously. Many are conditional: they apply only if certain criteria are met and cease to apply if circumstances change. An open-source model that becomes high-risk due to its deployment context loses its exemption. A research system deployed operationally is no longer covered by the R&D carve-out. The safest approach is to document why you believe an exclusion applies, including the facts supporting your conclusion, so you can demonstrate the basis for your position if challenged.

The next chapter maps the overall regulatory architecture: how the risk tiers work together, who bears what obligations in the AI value chain, and which bodies govern enforcement across Europe.

Chapter 3: The Regulatory Architecture

Mapping the Obligations Landscape

The Four Risk Categories Explained

Actors in the AI Value Chain

Governance Bodies: EU AI Office, Board, and National Authorities

Chapter 4: Prohibited AI Practices

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

The Unacceptable Risk Category

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Manipulation, Exploitation, and Scoring Prohibitions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Law Enforcement Exceptions and Safeguards

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Screening Your AI Practices for Prohibited Uses

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 5: High-Risk AI Systems: Qualification and Requirements

Annex I: The Product-Safety Route to High Risk

Annex III: The Standalone High-Risk List

Fundamental Rights Impact Assessments

When a System Becomes High Risk (and When It Does Not)

Chapter 6: Obligations for High-Risk AI Systems: Technical Requirements

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Data Governance and Quality Requirements

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Technical Documentation and Record-Keeping

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Transparency and User Information

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Accuracy, Robustness, and Cybersecurity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Human Oversight Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 7: Conformity Assessment and Market Surveillance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Internal Control versus Third-Party Assessment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

The CE Marking Process for AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Post-Market Monitoring Obligations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Incident Reporting and Corrective Actions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 8: General-Purpose AI Models and Systemic Risk

Defining General-Purpose AI Under the Act

Core Obligations for GPAI Providers

Systemic-Risk Models and Their Enhanced Duties

Open-Source and Foundation Model Considerations

Chapter 9: Transparency Obligations for Non-High-Risk AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Human-AI Interaction Disclosure Rules

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Generative AI Content Labeling

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Emotion Recognition and Biometric Categorization Transparency

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Practical Implementation of Transparency Measures

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 10: Obligations for Deployers and Other Actors

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Deployer Duties Before, During, and After Use

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Importers, Distributors, and Authorized Representatives

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Employee Use of AI: Special Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Supply Chain and Vendor Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 11: Regulatory Sandboxes, Innovation Support, and Standards

National AI Regulatory Sandboxes

The European AI Sandbox Network

Harmonized Standards and Presumption of Conformity

Codes of Practice and Soft Law Instruments

Chapter 12: Enforcement, Penalties, and Remedies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

National Supervisory Authorities and Their Powers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Administrative Fines and Penalty Tiers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Market Surveillance and Corrective Measures

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Private Enforcement and Individual Rights

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 13: Interaction with Other EU Laws

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

The AI Act and the General Data Protection Regulation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Overlap with the Digital Services Act

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

The Data Act and Data Access Obligations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Cybersecurity, Product Safety, and Sector-Specific Laws

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 14: Building Your EU AI Act Compliance Program

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Phase One: Discovery and Inventory

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Phase Two: Classification and Gap Assessment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Phase Three: Remediation and Documentation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Phase Four: Operationalization and Monitoring

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 15: Governance Structures and Organizational Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Designing an AI Governance Framework

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Accountability Models and Responsibility Matrices

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Integrating with ERM, Privacy, and Security Programs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Audit Readiness and Continuous Improvement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.

Chapter 16: Case Studies and Complex Scenarios

A Financial Institution Deploying Credit-Score AI

A Manufacturer Integrating AI into Medical Devices

A SaaS Provider Offering Generative AI Tools

A Multinational Corporation with Complex AI Supply Chains

Conclusion: The Future of AI Governance and What Comes Next

Where We Stand Today

Pending Delegated Acts and Evolving Interpretation

Beyond Europe: Global Regulatory Trajectories

What Organizations Should Prioritize

References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theeuaiactacompleteguidefororganizations.