The EU AI Act – A Practical Guide for U.S. and Other Non-EU Companies

Stéphanie De Giovanni | Roland Steidle | Mikhail Usubyan

September 23, 2026

The EU AI Act is applying in stages. Article 50 transparency duties now apply, and the Commission can enforce the general-purpose AI model regime. The Digital Omnibus delayed much of the high-risk regime, but it did not give businesses a general extension. The practical question is which obligations apply to your products and uses today.

A U.S. AI company may have no EU office, no EU employees and no EU data center – and still fall within the EU AI Act when it places AI systems or models on the EU market or its systems produce outputs used in the EU. [1]

The Act is not becoming relevant for the first time. Prohibited-practice and AI-literacy provisions have applied since February 2, 2025, and general-purpose AI (GPAI) model obligations began applying on August 2, 2025, subject to transitional rules. August 2, 2026 brought another important milestone: Article 50 transparency duties became applicable, the Commission gained its GPAI enforcement powers, and national authorities can enforce provisions already in effect. [1],[2],[3],[4]

The enforcement architecture is divided. The European Commission, acting through the AI Office, enforces obligations for GPAI-model providers and specified AI systems within its statutory remit. National competent authorities enforce the rules for other AI systems, while the European Data Protection Supervisor is responsible for AI systems used by EU institutions. The AI Office can request documentation and information, obtain access to models for evaluation, require corrective or risk-mitigation measures and, in appropriate cases, seek restriction, withdrawal or recall of a model. [1],[4]

The Commission has also launched complaint and whistleblower channels for potential violations falling within the AI Office's remit, together with a dedicated channel through which downstream system providers can raise alleged violations by GPAI-model providers under Articles 53–55. [1],[4]

The maximum fines are substantial. For example, Article 5 violations can attract up to €35 million or 7% of worldwide annual turnover; specified operator breaches, including Article 50 violations, can attract up to €15 million or 3%. For undertakings, the higher ceiling generally applies. Article 99 applies the lower fixed or percentage ceiling to SMEs, including start-ups, and extends that treatment to small mid-cap companies for specified breaches, but not for their Article 5 violations. Separately, Article 101 permits Commission fines of up to €15 million or 3%, whichever is higher, for intentional or negligent GPAI-provider infringements within its scope. Actual penalties must be proportionate, and Member States establish penalties for other infringements. [1],[2]

It would nevertheless be inaccurate to say that the entire Act is fully applicable. The July 2026 Digital Omnibus revised the timetable shortly before the August enforcement milestone. The result is a split compliance landscape: some obligations require evidence of compliance today, while the high-risk regime requires a credible, funded implementation roadmap. [2],[3]

For general informational purposes only. This publication does not constitute legal advice.

The current timetable at a glance

Date	What changes	Practical consequence
2 August 2026	Article 50 transparency duties apply. The Commission can enforce the GPAI regime, and national authorities can enforce provisions already applicable, including prohibited practices and AI-literacy duties.	Review live products and internal uses now. Earlier application dates still matter; this is not the start of all AI Act obligations.
2 December 2026	New Article 5 prohibitions concerning specified non-consensual intimate or sexually explicit material and child sexual abuse material apply. The limited Article 50(2) transition for certain pre-existing generative systems ends.	Complete the prohibited-practice screen and any legacy machine-readable marking implementation.
2 August 2027	Providers of GPAI models placed on the market before 2 August 2025 generally must comply with the GPAI regime.	Confirm model lineage, first market-placement date and whether a later release is legally the same model.
2 December 2027	The high-risk requirements for systems classified under Article 6(2) and Annex III apply, subject to transitional rules.	Build classification, governance, logging, documentation and conformity-assessment readiness well before the deadline.
2 August 2028	The high-risk requirements for systems classified under Article 6(1) and Annex I apply, subject to sector-specific and transitional rules.	Assess the product and safety-component criteria, and coordinate with the applicable product-regulatory and conformity-assessment framework.

Why the Act can reach companies outside Europe

The Act's territorial reach is not limited to companies incorporated or physically present in the EU. It applies to providers that place AI systems or GPAI models on the Union market, or put AI systems into service in the EU, whether the provider is established in the EU or a third country. It also applies to non-EU providers and deployers where output produced by an AI system is used in the EU, and it separately regulates importers, distributors, certain product manufacturers and authorized representatives. [1]

Common arrangements that can create EU exposure include:

- a U.S. software company offering an AI-enabled product to customers in one or more EU Member States;
- a U.S. model developer making a GPAI model, downloadable weights or a model API available on the EU market;
- a U.S. parent company or service provider operating an AI system outside the EU but supplying scores, recommendations or other outputs to an EU affiliate or customer for use in hiring, monitoring, credit or other decisions in the EU; and
- an offshore service provider producing AI-generated content, classifications or recommendations that a European customer uses in its business.

The analysis is also role-specific. A provider develops, or has developed, an AI system or GPAI model and places it on the market under its name or trademark, or puts an AI system into service under that name or trademark. A deployer uses an AI system under its authority. A company that accesses a third-party model through an API may not be the provider of the underlying model, but it may still be the provider of the customer-facing AI system it develops and markets. Contract labels matter commercially, but they do not override the statutory facts. [1],[9]

When the relevant high-risk provisions apply, Article 25 can shift provider obligations to another participant in specified circumstances – for example, where that participant places its name or trademark on an existing high-risk system, substantially modifies a high-risk system in a way that leaves it high-risk, or changes the intended

purpose of an existing system so that it becomes high-risk. That rule should not be generalized to every low-risk or transparency-risk system. [1]

For most businesses, a useful starting point is a proportionate map of their AI systems, uses and legal roles, rather than a generic "responsible AI" policy in isolation. That need not mean creating a separate, exhaustive register from scratch. Existing IT, procurement and data-protection records can be adapted to identify material AI uses and the areas requiring closer review.

What businesses must comply with now

1. Screen products and internal uses for prohibited practices

Some companies still treat the AI Act as a regime concerned mainly with future high-risk systems. That misses Article 5. The prohibited-practice rules have applied since February 2, 2025 and are now within the active enforcement framework. [1],[3],[4]

The review should extend beyond products marketed as AI. Ordinary business software may acquire AI features through vendor updates, sometimes with limited options to disable them. Procurement and IT teams should ask what has been added, what data it accesses, how it is used and whether settings or contractual controls are needed.

The prohibited practices include, subject to detailed statutory conditions, certain AI systems that use subliminal, manipulative or deceptive techniques to materially impair informed decision-making and cause or are reasonably likely to cause significant harm; exploit vulnerabilities associated with age, disability or socioeconomic circumstances; engage in prohibited social scoring; create facial-recognition databases through untargeted scraping of internet or CCTV images; infer emotions in workplaces or educational institutions, subject to limited medical or safety exceptions; or use biometric categorization to infer specified sensitive characteristics. Article 5 also contains rules on individual criminal-risk assessment and law-enforcement biometric identification that may be relevant to specialized products. [1]

Not every persuasive interface, personalization feature or behavioral model is prohibited. Several tests require specific intent, effects, harm or context. The practical response is a documented product and use-case review rather than a broad ban on "behavioral AI." Particular attention should be paid to HR tools, biometrics, age-sensitive services, behavioral targeting and products marketed as detecting emotion, intent, honesty or trustworthiness.

From December 2, 2026, Article 5 will also apply new prohibitions concerning specified AI systems that generate or manipulate non-consensual intimate or sexually explicit material and child sexual abuse material. The amended provision contains detailed conditions addressing intended purpose, reasonably foreseeable and reproducible outcomes and the presence or absence of reasonable safeguards, so product teams should not rely on a shorthand label such as "nudification ban" in place of a legal analysis. [1],[2],[3]

2. Implement a practical AI-literacy program

Article 4 applies broadly to providers and deployers, not only to developers of frontier models or providers of high-risk systems. As amended by the Digital Omnibus, it requires providers and deployers to take measures supporting the development of AI literacy among staff and other persons dealing with the operation and use of AI systems on their behalf. The measures should reflect technical knowledge, experience, education and training, the context in which the systems are used, and the people or groups affected. [1],[2],[8]

The revised provision does not require an organization to guarantee that every individual has reached a prescribed level of expertise. The Act does not mandate an "AI officer," external certification, a single curriculum or a particular governance structure. It does, however, require measures appropriate to the

company's role and risk profile, and national authorities can supervise the obligation. [2],[4],[8] Enforcement expectations are likely to continue evolving as national authorities develop supervisory practice.

A defensible program is likely to be role-specific. Engineers and product managers may need instruction on prohibited features, classification, logging and documentation. HR teams need to understand employment-related restrictions and high-risk use cases. Marketing and communications personnel need to recognize Article 50 labeling duties. Procurement, legal and deal teams need to identify missing vendor information, contractual dependencies and authorized-representative requirements.

An internal AI-use policy should turn that training into operational rules: which tools are approved, which information may be entered in prompts, which tasks require human review and which uses are restricted. For example, permission to help draft a reference should not be treated as permission to score an employee or recommend dismissal. Employers should also consider consistent access to approved AI assistance and training for people performing comparable tasks. These are practical governance choices, not a prescribed AI Act policy template. [14]

The Act does not expressly prescribe a training record or evidence file. As a practical matter, however, companies should retain reasonable evidence of the training, guidance, policies and other measures they have implemented. A generic annual presentation may be part of the program, but it is unlikely to be persuasive by itself where the organization develops or uses materially different AI systems across product, HR, marketing and operations.

3. Operationalize Article 50 transparency duties

Article 50 creates distinct obligations for providers and deployers of specified AI systems. It is not one generic "AI disclosure" rule, and the party responsible for the technical marking may be different from the party responsible for the human-facing label. [1],[5],[6]

Providers of AI systems that interact directly with natural persons – such as chatbots, AI agents and avatars – must generally design the system so that people are informed they are interacting with AI. The notice must be clear and distinguishable from the start of the first interaction, unless the AI nature of the interaction would be obvious to a reasonably well-informed, circumspect and observant person. Background systems, machine-to-machine communications and systems without direct human interaction may fall outside this particular duty. [1],[5],[6] In practice, providers should be cautious when relying on the obviousness exception and should document their reasoning, particularly for customer-facing assistants and AI agents.

Providers of systems, including general-purpose AI systems, that generate synthetic text, audio, images or video must generally ensure that the output is marked in an effective, reliable, robust and interoperable machine-readable format, as far as technically feasible, and can be detected as AI-generated or manipulated. The Commission's final guidelines address exclusions and qualifications for standard editing, source code, machine-to-machine outputs, closed-loop industrial uses and narrow business-to-business or industrial contexts. Those are fact-specific boundaries, not a general B2B exemption. [1],[5],[6]

Deployers have separate disclosure duties. They must inform people exposed to emotion-recognition or biometric-categorization systems. They must clearly disclose AI-generated or manipulated images, audio or video constituting a deepfake no later than first exposure, ordinarily through a visible or audible label that can be perceived without special tools. Machine-readable metadata inserted by the provider does not, by itself, satisfy that human-facing obligation. Deployers must also clearly label AI-generated or manipulated text published for the purpose of informing the public on matters of public interest, unless the text has undergone substantive human review or editorial control and a natural or legal person holds editorial responsibility for its publication. [1],[5],[6]

Article 50 also contains qualifications for evidently artistic, creative, satirical, fictional and analogous works, for which disclosure must not hamper display or enjoyment, and narrowly framed law-enforcement exceptions. Applicable notices must meet accessibility requirements. These qualifications should be assessed against the specific use, rather than treated as a general exception for entertainment or business communications. [1],[5],[6]

Article 50 has applied since August 2, 2026. The Digital Omnibus provides only a limited transition until December 2, 2026 for the Article 50(2) machine-readable marking and detection duty as applied to certain systems already placed on the market before August 2. It is not a general grace period for chatbot notices, deepfake disclosures or public-interest text labels. Content generated before August 2 does not have to be labeled retroactively, although the Commission encourages voluntary labeling where feasible. [2],[3],[6]

The voluntary Code of Practice on Transparency of AI-generated Content provides an EU-wide framework for demonstrating compliance with the marking and labeling obligations in Article 50(2), (4) and (5). The Commission and the AI Board have assessed it as an adequate voluntary tool. Non-signatories remain free to use alternative measures, but they must be able to demonstrate that those measures are adequate. [6],[7]

4. Determine whether the company is a GPAI-model provider

Not every company building a generative-AI product is a provider of a general-purpose AI model. A startup that accesses a third-party model through an API will ordinarily be a downstream system provider or deployer rather than the provider of the underlying model. It may nevertheless have its own Article 50, product-governance and, later, high-risk obligations. Upstream model compliance does not make the downstream product compliant. [1],[9]

A company that develops, or has developed, a broadly capable model and places it on the EU market under its own name or trademark, or makes a sufficiently significant downstream modification to an existing GPAI model, may be a GPAI-model provider under the Act and the Commission's guidelines. The legal assessment should consider model capability, training compute, who performed or commissioned the relevant development, the nature of the modification and how the model is supplied or used in the EU. Branding a release as merely a "fine-tune" or "open-source model" is not determinative. [1],[9]

Under Article 53, GPAI-model providers must draw up and maintain technical documentation about the model and its training, testing and evaluation; provide specified information and documentation to downstream system providers; implement a policy to comply with EU copyright and related-rights law, including rights reservations; publish a sufficiently detailed summary of the content used for training using the Commission template; and cooperate with the Commission and national authorities. [1],[9],[10]

A non-EU GPAI-model provider generally must appoint an EU authorized representative before placing the model on the Union market. Qualifying free and open-source models may benefit from exemptions from certain technical-documentation, downstream-information and authorized-representative duties, but the exemption is conditional, does not remove the copyright-policy or training-content-summary obligations, and does not apply to GPAI models with systemic risk. [1],[9],[10]

Providers of GPAI models with systemic risk face additional obligations: Commission notification where required, state-of-the-art model evaluation and documented adversarial testing, assessment and mitigation of systemic risks at Union level, serious-incident tracking and reporting, and cybersecurity protections for the model and its physical infrastructure. [1],[9],[10]

The Commission can now enforce these requirements. GPAI models placed on the market on or after August 2, 2025 are already subject to the regime; providers of models placed on the market before that date generally have until August 2, 2027 to comply. Providers may use the GPAI Code of Practice to demonstrate compliance or present alternative adequate measures for Commission assessment. [1],[3],[4],[9],[10]

High-risk AI: delayed, but not a reason to pause

The Digital Omnibus provides additional implementation time, but it does not eliminate the high-risk regime. Annex III addresses uses involving biometrics, critical infrastructure, education, recruitment and workforce management, creditworthiness, life and health insurance, access to essential services, emergency response, law enforcement, migration and border control, and the administration of justice. Employment applications include systems used to target job advertising, filter applications, evaluate candidates, influence promotion or termination, allocate tasks based on personal characteristics or monitor performance. [1],[2]

France provides a useful illustration of why national rules also matter. For employers covered by the CSE provisions applicable to companies with at least 50 employees, Article L2312-38 of the Labor Code requires prior information of the works council (CSE) about recruitment-assistance methods and automated personnel-management systems, including changes to them. Tools permitting employee monitoring require prior information and consultation before the implementation decision. An employment-AI assessment should therefore address existing national labor law alongside the AI Act. [12]

Annex III is not an automatic high-risk classification in every case. Article 6(3) contains a limited exception where the system does not pose a significant risk of harm to health, safety or fundamental rights, including by not materially influencing decision-making, and satisfies at least one of the specified conditions. These concern narrow procedural or preparatory tasks, improving the result of a completed human activity, or detecting decision patterns or deviations without replacing or improperly influencing a human assessment, subject to proper human review. Annex III systems that perform profiling of natural persons remain high-risk. A provider relying on the exception must document its assessment before market placement or putting into service and comply with the applicable registration requirements. This is a narrow carve-out, not the default rule for Annex III systems. [1]

The Commission's high-risk classification guidelines remain in draft form. They may help illustrate the analysis, but companies should distinguish draft guidance from the binding statutory tests and monitor the final text. [11]

When the relevant rules apply, providers of high-risk systems will face a substantial compliance architecture: risk management; data and data-governance controls; technical documentation; recordkeeping and logs; instructions and transparency for deployers; human-oversight measures; accuracy, robustness and cybersecurity; a quality-management system; conformity assessment; an EU declaration of conformity and CE marking where applicable; EU database registration; post-market monitoring; corrective action; and serious-incident reporting. Non-EU providers will generally need an EU authorized representative. [1]

Deployers will have their own responsibilities, including following instructions, assigning appropriately trained and empowered persons to human oversight, controlling relevant input data, monitoring operation, retaining logs under their control for the required period, generally at least six months unless applicable law provides otherwise, escalating risks and serious incidents, informing affected workers and, for specified Annex III decisions, informing individuals that a high-risk system is being used. Certain public bodies, public-service providers and deployers of creditworthiness or life and health-insurance systems must also conduct fundamental-rights impact assessments. [1]

A fundamental-rights impact assessment under the AI Act is distinct from a data-protection impact assessment under the GDPR. Where relevant issues have already been assessed under the GDPR, the AI Act assessment should complement that work rather than simply assume that one assessment satisfies both regimes. [1],[13]

Article 111 contains transitional rules for certain systems already placed on the market or put into service. Their application depends on the system category, the relevant market-placement date, public-authority use and subsequent significant design changes. Companies expecting to rely on a legacy-system transition should

preserve evidence of first market placement, model and system lineage, release dates and design changes rather than treating "legacy" as a label that can be asserted later. [1],[2]

Many of these requirements cannot be added shortly before a deadline. Logging, data provenance, version control, human oversight, testing, vendor dependencies and post-market monitoring often require product and process changes. December 2027 and August 2028 are therefore implementation runways, not permission to defer classification and architecture decisions.

The AI Act rarely operates alone

In our cross-border work with AI developers, technology companies, investors and multinational groups, a recurring practical lesson is not to treat the AI Act as a self-contained labeling statute. The same product or workflow may also require analysis under EU data-protection, cybersecurity, product-safety, consumer, employment, copyright, international-trade and sector-specific rules. Which regimes apply depends on the product, data, users and markets, but the analyses should be coordinated.

An Article 50 label does not cure unlawful processing of personal data. A vendor contract does not change statutory provider status. A conclusion that a system is not high-risk does not eliminate consumer, employment, intellectual-property or product-regulatory obligations. Conversely, an integrated assessment can often reuse the same system inventory, data map, testing evidence, incident process and contractual controls across several regulatory workstreams. [1],[13]

For non-EU businesses, this integration is particularly important because the commercial owner, model provider, developer, EU distributor, customer and affected individuals may be in different jurisdictions. The legal analysis must be translated into product requirements and contracts that the relevant parties can actually perform.

Four questions founders, boards, investors and acquirers should ask

1. What AI systems and models does the company actually provide or use – and in what legal role?

Management should be able to explain what AI the company provides or uses, and how it is classified. For deployers, a practical first screen can group uses into prohibited or restricted uses, potentially high-risk uses, Article 50 transparency cases and other lower-risk uses. These categories can overlap. Providers should separately assess their system- and model-level responsibilities. The objective is a reliable view of exposure, not another inventory that no one maintains.

Reuse existing IT and procurement records where possible and collect more detail for uses that warrant it: the system or model and version; business owner; intended and actual purpose; jurisdictions and users; data, outputs and human review; and the company's legal role. For products or models relying on a transition, preserve market-placement dates and version history. A lightweight process for vendor updates and new use cases is often more useful than a one-time exhaustive spreadsheet.

A statement that the company "uses only reputable third-party models" is not a classification analysis. It does not answer whether the company is the provider of the customer-facing system, whether outputs are used in the EU, whether a downstream modification creates GPAI-provider exposure or whether an Annex III use can rely on the Article 6(3) filter.

2. Can the company produce evidence of compliance with the rules already enforceable?

Depending on the company's role and use cases, relevant evidence may include a prohibited-practices assessment; AI-literacy materials and reasonable implementation records; chatbot notices and screenshots; testing of machine-readable marks and detection functionality; deepfake and public-interest text procedures; model documentation; copyright policies and training-content summaries; authorized-representative mandates; and a record of reliance on a Code of Practice or alternative compliance measures.

Diligence should also examine complaints, whistleblower reports, regulator inquiries, serious incidents, internal escalations and customer claims about AI Act compliance. The existence of a policy is less important than whether the company can show how it is implemented in the relevant product or workflow.

3. Do contracts across the AI value chain make compliance possible?

Upstream agreements should address access to technical and model documentation, notification of model changes, output-marking functionality, incident reporting, cybersecurity, intellectual-property and training-data issues, regulatory cooperation, audit or verification rights where appropriate, and termination or migration rights. Customer and distribution agreements should allocate operational tasks such as user notices, deepfake labels, input-data controls, human oversight and incident escalation.

For developers and providers selling to EU businesses, legal documentation is only part of customer readiness. Four product and contracting choices deserve early attention:

Hosting and deployment. EU data-center or on-premises options may be important to particular customers. They are not blanket AI Act requirements, and EU hosting does not itself resolve every data-protection or international-transfer issue. The architecture should match the customer's data, security and regulatory requirements. [1],[13],[14]

Use of customer inputs. Customers should be able to understand and control whether prompts, uploaded documents and outputs are retained or used to train public or shared models. A no-training commitment should be supported by the applicable service configuration and upstream provider terms, not merely by a sales assurance. [14]

External sources. Internet search, retrieval and other output-enrichment functions should be configurable so customers can control which sources are used and assess the resulting confidentiality, intellectual-property and data-protection risks.

Agent permissions. AI assistants and agents should respect user and task permissions, with narrower access and approval gates where appropriate. Replicating existing user access is not always enough: an agent can search, aggregate or act on information at a scale the user could not practically achieve unaided. Permission design should address that additional exposure. [13]

Contractual allocation is commercially important, but it cannot rewrite the company's statutory role. The objective is to ensure that the party bearing an obligation has the information, technical functionality, authority and cooperation needed to perform it.

4. Which current or planned products could become high-risk, and what is the implementation and transaction plan?

A company offering or deploying AI for hiring, employee management, education, credit, insurance, biometrics or another Annex III use should have an identified owner, budget and timetable for classification, data governance, technical documentation, logging, human oversight, conformity assessment, registration and post-market monitoring, as applicable to its role. The plan should also address whether an Article 6(3) filter or Article 111 transition is being claimed and what evidence supports it.

In financings and acquisitions, a missing roadmap may represent both regulatory exposure and technical debt. Depending on materiality, the issue can affect valuation, representations and disclosure schedules, remediation covenants, closing conditions, indemnification, customer-consent strategy and post-closing integration budgets. The diligence should be proportionate, but “we will address the AI Act later” is no longer a satisfactory response.

Compliance is becoming part of transaction readiness

The EU AI Act is not a single deadline and not a single compliance exercise. Its phased implementation combines enforceable duties for prohibited practices, AI literacy, transparency and GPAI models with a more demanding high-risk regime on a revised timetable.

For non-EU companies, the most serious mistake would be to interpret the high-risk postponement as a postponement of the Act itself. A business selling AI products into Europe, supplying models for European use or producing outputs used there should already be able to explain its role, identify the provisions that apply today, produce evidence of compliance and show a credible roadmap for the obligations that come next.

That work is not only about avoiding penalties. It increasingly affects enterprise sales, vendor reviews, insurance, financing diligence and M&A. The companies best positioned for European growth and transactions will treat AI Act compliance not as a legal appendix, but as part of product governance, commercial credibility and transaction readiness.

For general informational purposes only. This publication does not constitute legal advice.

About the authors

The authors collaborate across the firm's U.S., French and German practices on cross-border AI, technology, transactions, compliance and related regulatory matters. Their combined work spans AI developers and users, emerging and multinational companies, investment funds, boards and strategic investors.

Stéphanie De Giovanni	Partner, Rimôn Electa Law, Paris. Stéphanie advises on EU and French regulatory and compliance matters, international trade, product liability and related cross-border disputes. She specializes in the design and implementation of compliance programs.
Roland Steidle	Partner, Rimon Falkenfort, Frankfurt. Roland focuses on information technology, artificial intelligence, data protection and cybersecurity, and regularly publishes and teaches on technology and data-protection law.
Mikhail Usubyan	Partner, Rimon P.C., Washington, D.C. Mikhail advises private and public companies and investment funds on cross-border M&A, private equity, venture capital, strategic investments, technology matters and U.S.–Europe expansion.

* Stéphanie De Giovanni is a member of the French separate legal entity Rimôn Electa Law S.E.L.A.S., not Rimon P.C. Roland Steidle is a member of the separate German entity Rimon Falkenfort, not Rimon P.C.

Selected official sources

Law and guidance reviewed as of September 23, 2026. Commission guidance is interpretive, codes of practice are voluntary, and the high-risk classification guidelines listed at [11] remain a consultation draft. These materials do not replace the binding legislation. The practical recommendations in this article should be tailored to the business and its use cases.

[1] Regulation (EU) 2024/1689 (Artificial Intelligence Act), consolidated text as of 27 July 2026

[2] Regulation (EU) 2026/1744 (Digital Omnibus on AI)

[3] European Commission, Timeline for the Implementation of the EU AI Act

[4] European Commission, The enforcement framework of the AI Act

[5] European Commission, Guidelines on transparency obligations for providers and deployers of AI systems (final, 20 July 2026)

[6] European Commission, Transparency obligations under Article 50 of the AI Act – Questions and Answers

[7] European Commission, Code of Practice on Transparency of AI-generated Content

[8] European Commission, AI Literacy – Questions and Answers

[9] European Commission, Guidelines for providers of general-purpose AI models

[10] European Commission, General-purpose AI obligations under the AI Act

[11] European Commission, Draft guidelines on the classification of high-risk AI systems (19 May 2026; consultation draft)

[12] French Labor Code, Article L2312-38: recruitment methods, personnel management and employee-monitoring tools

[13] Regulation (EU) 2016/679 (GDPR), particularly Articles 5, 25, 28, 32 and 35, and Chapter V

[14] CNIL, How to deploy generative AI: initial guidance for organizations (18 July 2024; in French)

For general informational purposes only. This publication does not constitute legal advice.