FREE RESOURCE · FOR PRODUCT, ENGINEERING & COMPLIANCE LEADS

EU AI Act Readiness Checklist

A practical, eight-step audit you can run before legal asks. Built for teams who already have AI in production, or are about to.

Edition: May 2026 · Reflects Regulation (EU) 2024/1689 (the AI Act), applicable in phases from August 2024 to August 2027. The 2 August 2026 milestone (Annex III high-risk obligations) is roughly ten weeks out at this edition's date.

zyvra.studio · hello@zyvra.studio

START HERE

How to use this checklist

This is a working document. Each step is short on theory and long on questions you can actually answer. Print it, fill it in, and bring the gaps to your next product, legal, or engineering review.

Eight steps. Steps 1 to 3 decide whether the AI Act applies to you and how. Steps 4 to 6 are the obligations themselves. Steps 7 to 8 are timing and consequences. If you can tick every box in the steps that apply to you, you are in a defensible position. If you can't, you've found your roadmap.

Scope. The Act covers AI placed on the EU market, put into service in the EU, or where the output is used in the EU. Where you are headquartered does not matter. A US, UK, or Indian company shipping AI used by an EU customer is in scope.

Disclaimer. This checklist is informational and does not constitute legal advice. It distils the obligations of Regulation (EU) 2024/1689 into a working format for product and engineering teams; it is not a substitute for a qualified legal opinion. If your AI sits in healthcare, recruitment, financial services, education, essential public services, biometrics, law enforcement, migration, or justice, you should expect to need formal counsel.

> **Two-minute version.** If your AI makes or supports decisions about people, assume you are in scope, assume you are at minimum high-risk-adjacent, and assume you need automatic logging, human oversight, technical documentation, and a transparency notice. Steps 4 to 6 are how you prove it.

> **What's new in this edition (May 2026).** The 2 Aug 2026 milestone, when most of the Act including Annex III high-risk obligations becomes applicable, is now roughly ten weeks out. If you land in Tier 2, this edition is the last quarterly window before those obligations are enforced.

STEP 1

Are you in scope?

Tick everything that is true of your AI system. Any single tick keeps you in scope.

- [ ] The system is offered, sold, or made available to users located in the EU.
- [ ] The system is used inside the EU by an organisation we operate or supply.
- [ ] The output of the system (a decision, a score, generated content, a recommendation) is used in the EU, even if the model itself runs elsewhere.
- [ ] The system is used by an EU public authority, directly or via a vendor.
- [ ] We deploy the system as part of our own internal operations in an EU member state.

Out of scope (a partial list).

- [ ] The system is used exclusively for personal, non-professional activity by individuals.
- [ ] The system is part of an AI model still in pre-market research, testing, or development.
- [ ] The system is used solely for military, defence, or national security purposes.
- [ ] The system is released under a free and open-source licence, but only outside the high-risk and prohibited categories, and not for general-purpose AI models with systemic risk.

> If any box in the first list is ticked, continue to Step 2. If only the out-of-scope conditions apply, document why and revisit annually. Deployments expand.

STEP 2

Classify the risk tier

The Act sorts AI systems into four tiers. Find your highest applicable tier. You inherit every obligation at and below it.

Tier 1, Prohibited (Article 5) BANNED SINCE 2 FEB 2025

If you tick any of these, stop.

- [ ] Subliminal, manipulative, or deceptive techniques that materially distort behaviour.
- [ ] Exploitation of vulnerabilities tied to age, disability, or socio-economic situation.
- [ ] Social scoring of natural persons by public or private actors.
- [ ] Predictive policing based solely on profiling of an individual.
- [ ] Untargeted scraping of facial images from the internet or CCTV to build facial-recognition databases.
- [ ] Emotion recognition in workplaces or educational institutions (narrow exceptions).
- [ ] Biometric categorisation inferring race, political opinions, trade-union membership, religion, sex life, or sexual orientation.
- [ ] Real-time remote biometric identification in public spaces by law enforcement (narrow exceptions).

Tier 2, High-risk (Article 6 + Annex III) APPLIES FROM 2 AUG 2026

The bulk of the Act's obligations. Tick what applies.

- [ ] Safety component of a regulated product (machinery, medical devices, toys, vehicles, lifts, etc.) requiring third-party conformity assessment.
- [ ] Biometric identification, categorisation, or emotion recognition (outside Tier 1 bans).
- [ ] Critical infrastructure: road traffic, water, gas, heating, electricity, digital infrastructure.
- [ ] Education and vocational training: admission, evaluation, proctoring, prohibited-behaviour detection.
- [ ] Employment, worker management, access to self-employment: recruitment, screening, task allocation, performance and behaviour monitoring, promotion, termination.
- [ ] Access to essential services: public assistance, healthcare eligibility, credit scoring, insurance pricing for life and health, emergency call dispatching.
- [ ] Law enforcement: risk assessment of natural persons, polygraph use, evidence evaluation, profiling for criminal investigation.
- [ ] Migration, asylum, border control: risk and security assessments, visa decisions.
- [ ] Administration of justice and democratic processes.

Tier 3, Limited risk (Article 50) TRANSPARENCY OBLIGATIONS

- [ ] AI system that interacts with humans (chatbot, voice assistant).
- [ ] AI system that generates synthetic audio, image, video, or text content (generative AI output).

- [ ] AI system that performs emotion recognition or biometric categorisation (outside Tier 1 bans).
- [ ] AI system used to generate or manipulate deepfake content.

Tier 4, Minimal risk

Everything else. No mandatory obligations under the Act, but voluntary codes of conduct apply and good practice still matters.

STEP 3

Identify your role

The Act assigns different obligations to different actors. You can be more than one. A deployer that materially modifies a system, retrains it, or rebrands it can become a provider, and inherits provider obligations.

Provider

You develop the AI system or have it developed, and place it on the EU market or put it into service under your name or trademark.

- [ ] We develop the model and ship it to customers under our brand.
- [ ] We commission a model and deploy it as our own product.
- [ ] We materially modify a third-party AI system (retraining, fine-tuning, repurposing for a high-risk use).
- [ ] We re-brand or white-label a third-party AI system.

Deployer

You use an AI system under your authority, in your operations, for your customers, for your internal users.

- [ ] We use a third-party AI tool inside the company.
- [ ] We integrate a vendor AI into a product we offer to others.
- [ ] We rely on AI for HR, credit, eligibility, fraud, or other decisions about people.

Distributor, Importer, Authorised representative

Lighter obligations focused on verifying the provider has done their job.

- [ ] We resell or distribute someone else's AI system inside the EU.
- [ ] We import an AI system from outside the EU and place it on the EU market.
- [ ] We act as the EU representative for a non-EU provider.

> **Watch this trap.** Buying a foundation model API and wrapping it in a high-risk product (HR screening, credit scoring) makes you the provider of the high-risk system, regardless of who trained the underlying model.

STEP 4

High-risk obligations (Articles 8 to 15)

If Step 2 lands you in Tier 2, this is your obligations sheet. These are the controls regulators and enterprise buyers will ask to see.

Risk management system (Art. 9)

- [ ] Documented, continuous risk-management process across the full lifecycle.
- [ ] Identified known and reasonably foreseeable risks; assessed likelihood and severity.
- [ ] Risk-mitigation measures applied; residual risks judged acceptable and disclosed.
- [ ] Tested under realistic conditions before deployment, and re-tested after material change.

Data and data governance (Art. 10)

- [ ] Training, validation, and testing datasets are relevant, representative, and as far as possible free of errors and complete.
- [ ] Documented data sources, collection, labelling, cleaning, and enrichment.
- [ ] Examined for possible biases that may affect health, safety, or fundamental rights.
- [ ] Special-category personal data processed only where strictly necessary, with safeguards.

Technical documentation (Art. 11 + Annex IV)

- [ ] Documentation drawn up before market placement and kept current.
- [ ] Includes intended purpose, architecture, training data summary, performance metrics, limitations, and risk-management decisions.
- [ ] Available to national competent authorities on request, in a usable format.

Record-keeping, automatic logging (Art. 12)

- [ ] Automatic logging of events over the system's lifetime, enabled by design.
- [ ] Logs include period of use, reference databases checked, input data, and identification of natural persons involved in verification of results.
- [ ] Log retention period defined and consistent with intended purpose and legal duties.
- [ ] Logs are tamper-evident and exportable for audit.

STEP 4 (CONTINUED)

High-risk obligations, continued

Transparency to deployers (Art. 13)

- [ ] Instructions for use are concise, complete, accurate, and clear to deployers.
- [ ] Disclose intended purpose, performance characteristics, known limitations, human-oversight measures, and required computational and hardware resources.
- [ ] Disclose foreseeable misuse and the circumstances under which the system may behave unexpectedly.

Human oversight (Art. 14)

- [ ] Designed so that natural persons can effectively oversee the system in operation.
- [ ] Oversight measures appropriate to the risk and the autonomy of the system.
- [ ] Operators are able to fully understand the system's capacities and limitations and interpret its output correctly.
- [ ] Operators can decide not to use the system or to override, reverse, or stop its output.
- [ ] For Annex III biometric identification: at least two trained persons must verify a match before action is taken.

Accuracy, robustness, cybersecurity (Art. 15)

- [ ] Performance levels appropriate for intended purpose, declared in instructions for use.
- [ ] Resilience against errors, faults, and inconsistencies; fail-safe and redundancy where needed.
- [ ] Resilience against attempts to alter use, output, or performance through adversarial input, data poisoning, model evasion, or confidentiality attacks.
- [ ] Security controls proportionate to the risks and circumstances.

Provider obligations for high-risk systems (Arts. 16 to 22)

- [ ] Quality-management system in place, documented, and maintained.
- [ ] Conformity assessment completed before market placement (internal control or third-party, depending on Annex III category).
- [ ] EU declaration of conformity signed; CE marking affixed.
- [ ] System registered in the EU database for high-risk AI systems before market placement.
- [ ] Post-market monitoring system in place to collect, document, and analyse performance.
- [ ] Serious incidents reported to the relevant national authority within 15 days (72 hours for widespread infringements; 2 days for cybersecurity breaches affecting critical infrastructure).
- [ ] Corrective action taken when non-conformity is identified; affected parties informed.
- [ ] Authorised representative appointed if the provider is established outside the EU.

STEP 4 (CONTINUED)

Deployer specifics and FRIA

Deployer obligations for high-risk systems (Arts. 26 to 27)

- [ ] Use the system in accordance with the provider's instructions.
- [ ] Assign human oversight to natural persons with the necessary competence, training, and authority.
- [ ] Ensure input data is relevant and sufficiently representative for the intended purpose.
- [ ] Monitor operation and inform the provider of risks, serious incidents, or malfunctions.
- [ ] Keep system-generated logs for an appropriate period (minimum six months unless longer applies under EU or national law).
- [ ] Inform affected workers and their representatives before deploying a high-risk AI system in the workplace.
- [ ] Inform natural persons subject to a decision produced or assisted by the system, where the decision produces legal or similarly significant effects.
- [ ] Cooperate with national authorities on any action taken in relation to the system.

Fundamental Rights Impact Assessment (Art. 27)

Required for deployers that are public bodies, private operators providing public services, or operators using high-risk systems in creditworthiness, life and health insurance pricing, biometric identification, or emotion recognition. Must be carried out before first deployment.

- [ ] Description of the deployer's processes in which the high-risk AI will be used.
- [ ] Period of time and frequency of intended use.
- [ ] Categories of natural persons and groups likely to be affected.
- [ ] Specific risks of harm to those persons, taking into account information from the provider.
- [ ] Description of the human-oversight measures in place.
- [ ] Measures to be taken if those risks materialise, including governance and complaints handling.

STEP 5

Transparency obligations (Article 50)

These apply on top of any high-risk obligations, and they apply regardless of risk tier whenever you fit the description. Most consumer AI products tick at least one.

- [ ] Users interacting with an AI system are informed they are interacting with AI, unless this is obvious to a reasonably well-informed person in context.
- [ ] Synthetic audio, image, video, or text content is marked in a machine-readable format detectable as artificially generated or manipulated.
- [ ] Deepfakes, content that appreciably resembles real people, objects, places, or events, are clearly labelled as artificially generated or manipulated.
- [ ] Text generated or substantially modified by AI and published to inform the public on matters of public interest is disclosed as such (with editorial-control exceptions).
- [ ] Emotion recognition or biometric categorisation systems inform the natural persons exposed to them about the operation of the system.

> **Watermarking is now table stakes.** If you ship a generative model or use one to create customer-facing content, you need a documented technical method for marking your output as AI-generated. Plan the watermarking story before launch. Retrofitting a content-provenance pipeline is painful.

STEP 6

General-purpose AI model obligations (Articles 51 to 55)

Apply if you train or substantially modify a general-purpose AI model (a foundation model). Most product teams using a third-party API are not providers of GPAI, but if you fine-tune at scale, this matters. **These obligations have applied since 2 August 2025.**

All providers of GPAI models

- [ ] Technical documentation for the model maintained and made available to the AI Office and national authorities on request.
- [ ] Documentation made available to downstream providers integrating the model.
- [ ] Policy in place to comply with EU copyright law, including respecting opt-outs from text and data mining.
- [ ] Sufficiently detailed summary of training data published, using the AI Office template.

GPAI with systemic risk

Training compute above 10^{25} FLOPs, or designated by the Commission.

- [ ] Notify the Commission within two weeks of meeting (or being expected to meet) the threshold.
- [ ] Model evaluation including adversarial testing to identify and mitigate systemic risks.
- [ ] Track, document, and report serious incidents and possible corrective measures.
- [ ] Adequate cybersecurity protection for the model and its physical infrastructure.

STEP 7

Timeline, what's already on you and what's coming

The Act entered into force on 1 August 2024 and applies in phases. Plan against your latest applicable date, not the earliest.

DATE	WHAT APPLIES	STATUS, MAY 2026
1 Aug 2024	AI Act enters into force.	Done
2 Feb 2025	Prohibitions (Art. 5) and AI literacy obligations (Art. 4) apply. Anyone whose staff use AI on the job needs an AI-literacy programme.	Live
2 Aug 2025	GPAI model rules (Arts. 51 to 55), governance, notifying authorities, and most penalty provisions apply. National competent authorities must be designated.	Live
2 Aug 2026	**Most of the Act applies, including Annex III high-risk systems.**	**≈10 weeks**
2 Aug 2027	Annex I high-risk systems (safety components of regulated products) and the long-tail conformity-assessment provisions apply.	Coming

> **What this means today.** If your build lands in Tier 2 (high-risk under Annex III), the 2 Aug 2026 milestone is the line. Article 12 logging and Article 14 human-oversight measures are the two items that must be designed into the product, not bolted on after the fact.

STEP 8

Penalties, what non-compliance actually costs

Fines are tiered and turnover-linked, with the higher of a fixed amount or a percentage of worldwide annual turnover applying.

BREACH	MAXIMUM FINE
Use of prohibited AI (Art. 5)	€35 million or 7% of global annual turnover, whichever is higher.
Non-compliance with most other obligations (data, transparency, high-risk requirements, deployer duties)	€15 million or 3% of global annual turnover, whichever is higher.
Supply of incorrect, incomplete, or misleading information to authorities	€7.5 million or 1% of global annual turnover, whichever is higher.
For SMEs and start-ups	The lower of the two amounts applies (capped to protect smaller actors), but the obligations are the same.

Beyond fines. A national authority can require withdrawal from the EU market, order corrective action, and publicise the decision. For B2B vendors, the harder cost is usually the enterprise procurement door closing. Large EU buyers are now refusing AI vendors that cannot evidence Article 12 logging and transparency obligations.

NEXT

What to do with this checklist

If you ticked nothing in Step 1

You are likely out of scope today. Re-run this checklist when you onboard EU customers, EU users, or EU staff, and at least once a year regardless. Scope creeps quietly.

If you landed in Tier 3 (limited risk)

Your priority is Step 5, transparency. Implement an AI-disclosure pattern across your interfaces and a watermarking method for any synthetic content you generate. Document both. The work is small if you do it now.

If you landed in Tier 2 (high-risk)

Treat the gaps in Step 4 as a roadmap. The hardest items to retrofit are Article 12 logging (must be designed in) and Article 14 human-oversight hooks (must be in the product surface, not bolted on). Start there, then work through documentation and the quality-management system in parallel. The 2 Aug 2026 milestone is your hard deadline.

If you landed in Tier 1 (prohibited)

Stop the deployment. Review with counsel before re-scoping. Several Tier 1 categories have narrow exceptions, but the burden of proof sits squarely on you.

A SECOND PAIR OF EYES

Want help closing the gaps?

Zyvra Studio Ltd builds security-first software at startup speed. Architected to enterprise standards, shipped in weeks, signed off by your security team. Products may or may not include AI capability. When they do, scoping folds in the EU AI Act dimensions alongside the universal ones.

If you'd like a 30-minute readiness review against this checklist (no fee, no pitch), email **hello@zyvra.studio** with the words *readiness review* in the subject line. We'll come back within one business day.

The work is one fixed-price **Build & Ship** engagement (scope, build, ship), priced by tier, with introductory launch rates locked in for your engagement (rising after 30 September 2026): zyvra.studio/services/build-and-ship.html

Essential, £2,000	A single-task build, end to end. 1 to 2 integrations, single cloud region, not in a regulated sector. Includes scoping (recommended architecture and 3-month cost model), the build deployed in your own environment, and full code, IaC, and runbook handover with 60 days of support.
Standard, £5,000	A multi-step workflow across two or three production systems, with an audit-logging baseline and some compliance considerations. Everything in Essential, plus cost tracked against the model and custom auth or network policies where needed.
Regulated, £10,000	A regulated-sector deployment (FCA, MHRA, ICO) with Article 12 compatible decision logging from day one, a full evidence pack for your compliance team, and support for an external auditor review. Your team owns the system the day we ship.