EU AI Act readiness checklist

EXPERANOS Research & Innovation | a practical self-assessment for AI system owners

The EU AI Act phases in between 2025 and 2027. If your organisation develops or deploys an AI system in the EU, walk through this checklist. Every unticked box is a work item, and most of them are engineering, not paperwork. This is an orientation aid, not legal advice.

1. Scope and risk classification

- [ ] We have an inventory of every AI system we develop, procure or operate, including 'hidden' ones inside vendor products.
- [ ] Each system is classified: prohibited practice / high-risk (Annex III or safety component) / limited risk (transparency duties) / minimal risk.
- [ ] For high-risk candidates, the classification rationale is written down and has been reviewed by someone accountable.
- [ ] General-purpose AI models we build on (LLMs, foundation models) are identified, with their providers' documentation on file.

2. Data and governance

- [ ] Training, validation and test datasets are documented: origin, licensing, collection method, known gaps and biases.
- [ ] Data relevance and representativeness for the intended purpose have been examined and recorded.
- [ ] Personal data use has a GDPR legal basis, and a DPIA exists where processing is high-risk.
- [ ] Dataset versions are tracked so any model can be traced back to the exact data that produced it.

3. Technical documentation and record keeping

- [ ] Each high-risk system has technical documentation: intended purpose, architecture, training process, performance metrics and limitations.
- [ ] Model cards and datasheets exist and are kept current with each release.
- [ ] The system automatically logs events relevant to risk (inputs, decisions, overrides, failures) with defined retention.
- [ ] Logs are tamper-evident or otherwise protected, and someone is assigned to review them.

4. Accuracy, robustness and cybersecurity

- [ ] Accuracy metrics are defined for the intended use, measured on realistic data, and stated honestly in the documentation.
- [ ] Robustness has been tested: distribution shift, edge cases, degraded inputs and, where relevant, adversarial attacks.
- [ ] Bias and fairness have been evaluated across the population groups the system affects, with results recorded.
- [ ] The AI pipeline is covered by cybersecurity controls: access management, dependency scanning, model and data integrity checks.

5. Human oversight and transparency

- [ ] A named role can inspect, override or stop the system, and the interface genuinely enables that in the time available.
- [ ] Operators are trained on the system's limitations and the situations in which it should not be trusted.
- [ ] People interacting with the system are informed they are dealing with AI where the Act requires it (chatbots, emotion recognition, deepfakes).
- [ ] Instructions for use exist for deployers, including foreseeable misuse.

6. Post-market monitoring and incidents

- [ ] Model performance is monitored in production: drift, error rates and unexpected behaviour raise alerts to a defined owner.
- [ ] A post-market monitoring plan exists and feeds real production data back into risk management.
- [ ] There is an incident process that can meet the Act's serious-incident reporting deadlines.
- [ ] Substantial modifications trigger re-assessment before deployment, not after.

7. Organisation and timeline

- [ ] Accountability for AI Act compliance is assigned to a named person or team with budget.
- [ ] The phase-in dates relevant to your systems are mapped against your release calendar.
- [ ] Suppliers and integration partners have been asked for their conformity evidence.