The EU AI Act Product Quick-Check

Place your AI features under the EU AI Act before the high-risk obligations bind on 2 August 2026 — the four risk tiers, the Annex III triggers, and the five things a product team must do now.

This is written for product managers, not lawyers. The EU AI Act is not a future hypothetical: the obligations for high-risk systems bind on **2 August 2026**. If your product touches the EU market and uses AI, the first job is not compliance — it's classification. You cannot scope the work until you know which tier you're in.

Run each AI feature through this quick-check before the next sprint plans around it.

Step 1 — Place each feature in a risk tier

The Act sorts AI systems into four tiers. Most product features land in the bottom two; the cost of being wrong lives in the top two.

- **Unacceptable risk — prohibited.** Social scoring, manipulative or exploitative systems, most real-time biometric identification in public spaces. If a feature is here, it doesn't ship in the EU. Full stop.
- **High risk — heavy obligations.** Systems used in the contexts listed in Annex III (see Step 2), or that act as a safety component of a regulated product. This is where the 2 August 2026 duties land.
- **Limited risk — transparency obligations.** Chatbots, generative content, emotion recognition. Users must be told they're interacting with AI; synthetic media must be labelled.
- **Minimal risk — no specific obligations.** Spam filters, recommendations, most internal tooling. Document the call and move on.

Step 2 — Check the Annex III triggers

A feature is high-risk if it's used in (among others):

- Biometrics and biometric categorisation.
- Critical infrastructure — energy, water, transport — as a safety component.
- Education and vocational training: access, scoring, proctoring.
- Employment: recruitment, screening, task allocation, evaluation.

- Access to essential services: credit scoring, insurance pricing, public benefits.
- Law enforcement, migration, and the administration of justice.

If your feature decides, ranks, or gates a person's access to one of these, assume high-risk until proven otherwise.

Step 3 — The five things a product team must do now

Whatever the tier, these five move you from exposed to defensible:

1. **Inventory.** List every AI system and feature you ship or deploy. You can't classify what you haven't named — and the inventory is the single most urgent step.
2. **Classify.** Assign each one a tier and write down why. This document is your first line of defence.
3. **Assign an owner.** One accountable person, or an AI committee with a named chair, for AI governance — not a shared inbox.
4. **Transparency.** For anything user-facing, tell people they're dealing with AI and label generated content. Cheap to do, expensive to skip.
5. **Human oversight.** For high-risk systems, build in the ability for a person to understand, override, and detect anomalies — in the architecture, not the release notes.

One caveat

The Act has no finish line. The Commission is still issuing guidance, the AI Office is revising codes of practice, and national regulators are operationalising penalties. Treat classification as a living document you revisit each quarter, not a one-time gate.

Getting the classification right early is the cheap part; discovering you were high-risk after launch is the expensive part. If you're staring at a backlog of features and aren't sure which ones bind on 2 August, that triage is exactly the kind of thing I help product teams work through.