THE EU AI ACT BUILDER'S CHECKLIST

27 checks across inventory, risk classification, documentation, incident reporting and procurement — written for the people who ship and run AI systems, not for law firms. By Kris Kassem · flamedfox.online · The Program Recovery (weekly letter) · v1.0 — September 2026

> **Why now:** the EU AI Act's high-risk obligations became applicable on **2 August 2026**. Under the AI Omnibus adjustment, a limited grace period for certain legacy systems runs to **2 December 2026**. If your organization ships or operates AI that touches the EU — as a provider or a deployer — the window to get your house in order is now, and it is measured in weeks. **What this is:** the working checklist I use with program sponsors to scope AI Act readiness. It is practical, not legal advice — pair it with counsel for binding interpretations.

HOW TO USE THIS CHECKLIST

Score each check: ✅ done · 🟡 partial · ❌ not started · ➖ not applicable (write down why). Anything ❌ in Sections A–B blocks everything downstream. Do those first. 60–90 minutes for a first pass with your delivery lead; the output is your readiness gap list.

A. INVENTORY (do this before anything else)

1. ☐ You have a complete AI system inventory. Every system that uses AI in a product or internal process — including vendor products with AI features, embedded models, and "just a script" automations. Owners named for each entry. **2. ☐ Each system has a provider/deployer label.** Under the Act your duties differ radically depending on whether you *provide* a system (you built/brand it) or *deploy* it (you use it under your own authority). Many organizations are both, per system. **3. ☐ Each system records its data flows.** What personal or sensitive data goes in, what comes out, where it is stored, which countries touch it. **4. ☐ AI features you bought are inventoried too.** Your ATS, CRM, support desk and office suite have shipped AI features this year. "We didn't build it" is not an exemption — deployer duties follow the use. **5. ☐ The inventory has an update mechanism.** An inventory nobody maintains is a liability with a timestamp. Name who updates it, when, and on what trigger (new system, vendor update, scope change).

B. RISK CLASSIFICATION

6. ☐ Each system is mapped against the prohibited practices. (e.g., manipulative techniques, social scoring, untargeted facial scraping, emotion inference in workplaces/education, certain biometric categorization.) If anything touches these: stop and escalate — the prohibition applies and penalties are the highest tier. **7. ☐ Each system is checked against the high-risk categories.** Employment/worker management, education, essential services (credit, insurance, benefits), law enforcement, migration, and AI as a safety component of regulated products are the big ones for most enterprises. **8. ☐ Embedders and GPAI wrappers are classified.** A system built on a general-purpose model inherits duties from how you use it. Document the reasoning, not just the label. **9. ☐ Borderline calls are documented.** For every "we decided this is limited-risk/minimal-risk," write down the argument and who signed it. Regulators and auditors ask for the reasoning, not the conclusion.

C. IF HIGH-RISK — THE SYSTEM BUILD-OUT

10. ☐ **A risk management process exists and is running.** Continuous, documented, proportionate: identify → evaluate → mitigate → residual risk accepted by a named person. **11.** ☐ **Data governance is documented per system.** Training/validation/testing data quality, bias examination, representativeness — to the depth appropriate to your role (provider vs deployer). **12.** ☐ **Technical documentation exists per system.** The Annex IV package for providers; deployers keep at minimum the vendor documentation, their configuration, and their use restrictions. **13.** ☐ **Logging is real, not aspirational.** Automatic event logs with the traceability the Act expects. "The model vendor probably logs something" is not an answer. **14.** ☐ **Users get the required information.** Deployers and end-users must know they are interacting with AI where applicable, what it does, and its limits. Check your UIs, release notes, and internal SOPs. **15.** ☐ **Human oversight is designed, not declared.** Named humans with real authority and competence to intervene, override or stop the system. "The manager can email the vendor" fails this. **16.** ☐ **Accuracy, robustness and cybersecurity are addressed.** Consistent performance targets, testing evidence, and security posture proportionate to the risk. **17.** ☐ **Accessibility is considered.** High-risk systems must be accessible for persons with disabilities to the extent feasible. Frequently forgotten; cheap to note now.

D. DEADLINES & TRANSITION

18. ☐ **You know which deadline applies to each system.** High-risk (Annex III) obligations apply from 2 Aug 2026; high-risk embedded in regulated products from 2 Aug 2027. Put them on the program calendar with owners. **19.** ☐ **Legacy-system grace is mapped to 2 Dec 2026.** If the Omnibus grace period applies to systems you placed on the market before 2 Aug 2026, treat 2 December 2026 as the hard wall: list what must be true by then. **20.** ☐ **The work is scheduled backwards from the wall.** Documentation and logging take longest. If they start after October, the plan is fiction. **21.** ☐ **A decision-maker owns AI Act readiness.** One accountable owner (not a committee). For most mid-market organizations this is the CIO/CTO or the AI program sponsor.

E. INCIDENT REPORTING & MONITORING

22. ☐ **The serious-incident reporting path is documented.** Providers must report serious incidents to market surveillance authorities; deployers have information duties. Who detects, who decides, who files, within what timeline? (The Act's regime includes short reporting windows — your process must be faster than your legal review cycle.) **23.** ☐ **Post-market monitoring is planned.** Provider-side monitoring per the Act's expectations; deployer-side equivalents for the systems you operate. **24.** ☐ **A "stop and reassess" trigger exists.** Define now what performance drift, misuse or incident pattern forces a reclassification or a pause. Written down before you need it.

F. VENDOR & PROCUREMENT

25. ☐ **Vendor AI contracts carry the right warranties.** Compliance representations, documentation delivery, incident notification duties, cooperation on investigations, audit rights, indemnities proportionate to risk. **26.** ☐ **Procurement asks the AI Act question by default.** "Which obligations does this product put on us as a deployer, and what documentation do you provide?" — added to every AI-touching purchase. **27.** ☐ **Third-party AI (APIs, embedded models) is on the inventory radar.** New vendor features ship silently. Your inventory process (Check 5) must catch them.

SCORING

- **0–5 ❌ in A–B:** you have a scope problem, not a compliance problem yet. Fix inventory and classification within 2 weeks.
- **❌ in C (high-risk duties):** the Dec 2 wall is real for you. Documentation and logging first — they gate everything else.
- **❌ in E–F:** you are one vendor incident away from finding out your reporting path doesn't exist.

WHAT I'D DO THIS WEEK (if it were my program)

1. Run Sections A–B with your delivery lead (60–90 min, this checklist as the agenda).
2. Assign the one accountable owner (Check 21) — in writing.
3. Start the technical documentation for the highest-risk system — it is always slower than anyone expects.
4. Add the procurement question (Check 26) to your next vendor review cycle.
5. Put 2 December 2026 on the program calendar with the specific legacy systems named.

Kris Kassem is an enterprise AI delivery and program-turnaround consultant (20+ years, $400M+ delivered in banking and resources; MBA, CISM, PMP, ITIL Expert). He writes The Program Recovery, a weekly letter for AI program sponsors: https://theprogramrecovery.substack.com · Templates and consulting: https://flamedfox.online