Article 50 applies now.

The EU AI Act's transparency duties apply from 2 August 2026. The main high-risk deadlines moved. **These did not.**

€15M

UP TO, OR 3% OF TURNOVER

Article 50 violations can be fined up to €15 million or, for an undertaking, **3% of total worldwide annual turnover** for the preceding financial year, whichever is higher.

Regulation (EU) 2024/1689 · Article 99(4)(g) · Article 50 applies from 2 August 2026

SMEs and small mid-cap companies are subject to lower-cap rules.

The 2026 amendment moved the main high-risk deadlines to 2 December 2027 and 2 August 2028. **Article 50 still applies from 2 August 2026.**

One narrow transition remains: providers of systems generating synthetic audio, image, video or text that were placed on the market before 2 August 2026 have until 2 December 2026 to comply with Article 50(2)'s machine-readable marking duty.

The important question now is not whether your organization has an AI policy. It is whether the people building, buying, deploying, governing, and operating these systems know which obligations belong to them.

Up to €35M / 7%

Prohibited AI practices under Article 5.

Not the transparency tier.

Up to €15M / 3%

Article 50 transparency obligations, plus specified operator and notified-body duties.

This is the relevant ceiling.

Up to €7.5M / 1%

Incorrect, incomplete or misleading information supplied in response to a request from a notified body or national competent authority.

If it talks to a person

Tell people they are interacting with AI unless that would be obvious to a reasonably informed, observant and circumspect person in context.

Provider duty. Chatbots, assistants, voice agents, intake tools.

If it generates content

Machine-mark synthetic audio, image, video or text so the output is detectable as artificially generated or manipulated.

Provider duty. Technical-feasibility and editing exceptions apply.

If it reads people or publishes

Tell people when they are exposed to emotion recognition or biometric categorisation. Disclose deepfakes and covered AI-generated or manipulated text published to inform the public on matters of public interest.

Deployer duty. The Act's exceptions apply; GDPR and other applicable law still sit alongside it.

THE PART NOBODY BUDGETED FOR

Regulation becomes an operating system. The law is written in articles. Your company runs through people.

Legal interprets it. Technology implements it. Procurement buys around it. Product makes tradeoffs. Operations discovers the edge cases. Someone eventually has to explain the result. Compliance fails in the gaps between those groups. **That is not primarily a legal drafting problem. It is an organizational one.**

Know what is actually in scope

Inventory systems that interact with people, generate synthetic content, perform emotion recognition or biometric categorisation, or publish covered material.

Not the systems listed in last year's AI strategy. The systems actually running.

Name the human who owns the duty

Provider and deployer obligations are different. The same technology can create different responsibilities depending on the organization's role. Every requirement therefore needs a person, not simply a policy.

Who owns the disclosure? Who verifies the marking? Who approves an exception? If everyone owns compliance, nobody does.

Leave evidence behind

For every in-scope system: owner, applicable Article 50 paragraph, disclosure or marking method, exception relied on, system version, and evidence that the control is operating.

A policy says what you intended to do. Evidence shows what happened. That distinction has occupied most of my professional life.

WHY I DO THIS DIFFERENTLY

I started my career in ruins. Regulators leave considerably better documentation. Anthropology taught me to look beyond the formal rule to the system of people interpreting it. Engineering taught me to keep following the evidence until I could show exactly what the technology did. Both matter here.

The EU AI Act is not going to be implemented by Legal alone, and it is not going to be implemented by Engineering alone. The durable answer is a common language between the people interpreting the regulation, the people building the system, the people operating it, and the executives accepting the risk.

That agreement becomes architecture. Architecture leaves evidence. And evidence is considerably easier to defend than confidence.

You do not need to wait for perfect regulation. The rules will continue to evolve. So will the technology. The useful standard is not predicting every future requirement. It is building an organization that knows what it intended, knows what it deployed, knows who owns the decision, and can reconstruct what happened. Do that well and regulatory change becomes manageable rather than existential. It also happens to be good engineering.

Article 50 exposure read: fixed scope, two weeks. Standing advisory is limited to three relationships because the useful work happens across legal, technical, product, and executive leadership.

This is not a compliance deck passed between departments, a legal opinion, or a rubber stamp. Where the underlying AI decision also needs to be explainable, Bast adds bounded retrieval, ontology grounding, and traceable source lineage. Operational and technical advisory; not legal advice.

Beth Rudden · Founder & CEO, Bast AI · beth@bast.ai · bethrudden.ai