GECIĆ | LAW

WHAT THE AI ACT DEAL REALLY MEANS

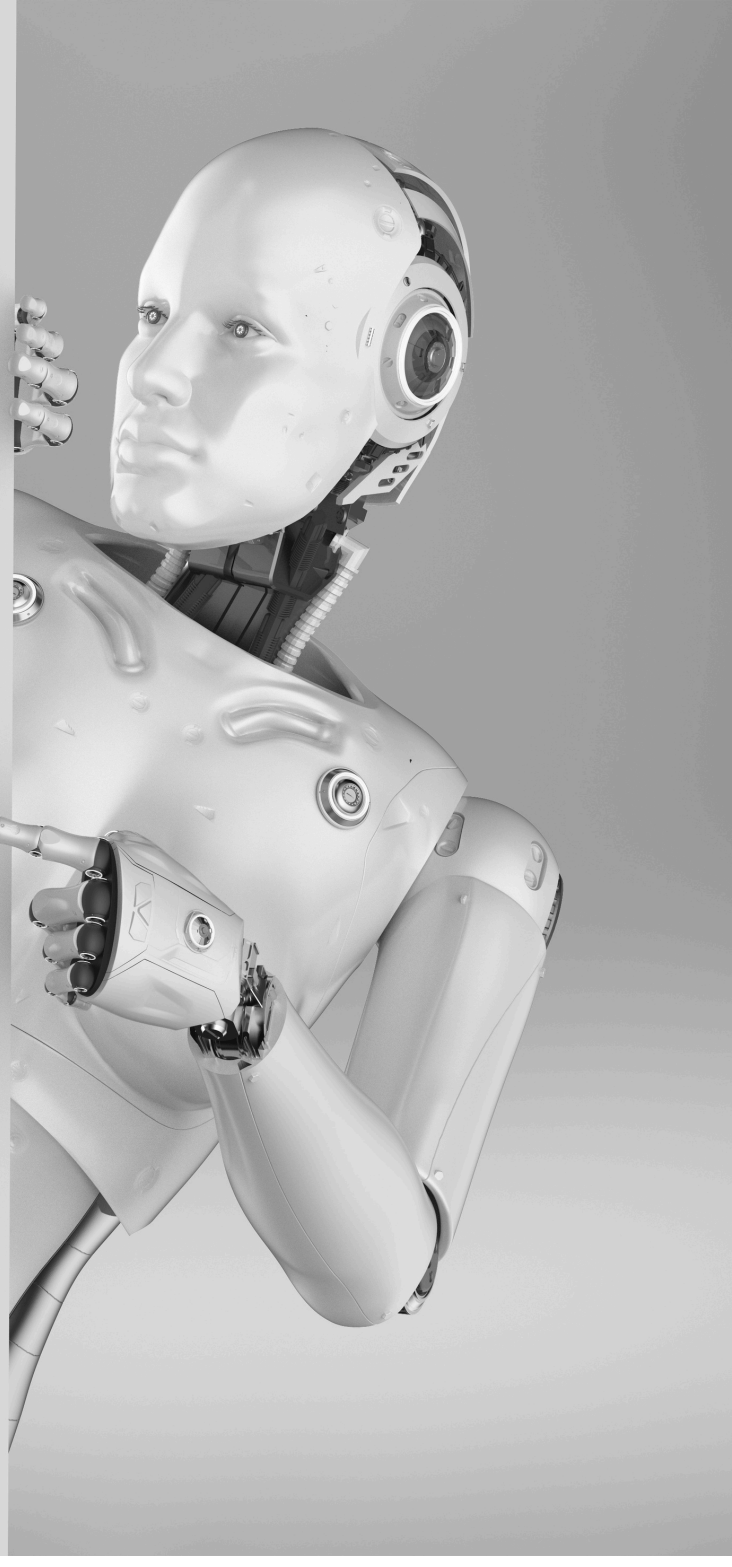

BY UROŠ RAJIĆ

The European Union reached a provisional agreement to simplify its flagship artificial intelligence regulation. Member states and the European Parliament concluded a deal to delay key obligations for high-risk AI systems and streamline compliance requirements for businesses.

"We now make the AI rules more workable in practice, remove overlaps and pause the high-risk requirements. In order for Europe to become an AI continent, we need to promote innovation, support startups and scaleups and make it easier to build AI in Europe."

That reads well. It also, rather quietly, sidesteps the harder question: if the requirements have not changed and nothing substantively new has been introduced, what exactly accelerates Europe's position in its race with the rest of the world.

What Has Changed, and What Has Not

The headline change is a timeline adjustment. High-risk AI systems covered under Annex III of the AI Act – those deployed in employment, education, credit scoring, and law enforcement – will now face compliance obligations from December 2, 2027, pushed back from the original August 2026 deadline. AI systems embedded in regulated products under Annex I – medical devices, machinery, toys – have until August 2, 2028. The provisional agreement removes overlapping AI Act obligations for machinery covered by sectoral product-safety legislation, so that AI-related safety requirements for machinery are addressed through the Machinery Regulation rather than through direct duplicate application of the AI Act.

It is worth being precise about what this extension changes, because it changes less than it might appear. The core obligations for deployers of high-risk AI systems, rooted in Article 26, remain intact. What the extension provides is additional time before national market surveillance authorities can enforce them. The obligations themselves (provider instructions, human oversight, serious incident reporting, etc.) are not suspended. They are deferred. The distinction matters, particularly for deployers who operate in sensitive sectors where reputational and liability exposure arises long before a regulator issues a notice.

Several other changes take practical effect sooner. AI systems that interact with users, generate synthetic

content, or perform emotion recognition must comply with transparency obligations under Article 50 from August 2, 2026. This date has not moved. This date remains unchanged for general disclosure obligations: chatbot identification, deepfake labelling, and emotion recognition disclosure. However, the provisional agreement shortened the watermarking grace period for AI-generated audio, image, video, and text from six months to three months, placing that specific sub-obligation on December 2, 2026.

New additions in today's deal include a ban on AI nudification applications capable of generating non-consensual intimate imagery, and a mandatory watermarking obligation for AI-generated audio, image, video, and text content, applicable from December 2, 2026. In addition, compliance relief previously available to SMEs is extended to small mid-caps, defined as companies with fewer than 750 employees and either annual turnover below €150 million or annual balance sheet below €129 million. The deal reinstates EU database registration for providers relying on the Annex III "no significant risk" derogation, rather than replacing registration with purely internal documentation.

The agreement is provisional. Formal adoption by Parliament and Council is still required. If that process is not complete before August 2, 2026, the original unamended text applies.

The Mechanism of Adaptation and Its Practical Limits

The Commission holds the power to amend the list of high-risk use cases in Annex III through delegated acts. A separate obligation requires periodic review of the Act's prohibited AI practices 5. The practical effect is the same: the high-risk list is not fixed permanently and can be updated by Commission initiative, subject to the procedural constraints described below.

In practice, adaptation is not quick. Under Article 97, delegated acts are subject to Member State expert consultation and notification to Parliament and Council. Parliament and Council have three months to object, extendable by a further three months. On an accelerated track, the best-case scenario is three to four months. Whether that is fast enough for a technology whose frontier moves quarterly is a reasonable thing to consider. The Digital Omnibus that produced today's deal was proposed in November 2025. The deal followed difficult trilogue negotiations, including a failed round on 28 April 2026 and a further overnight session that produced agreement in the early hours of 7 May 2026.

Whether the architecture of the Act, such as fixed categories, enumerated use cases, documentary compliance procedures, and a standardization process built for physical products, is well-matched to the speed at which AI applications move is a structural question the deal does not address.

Two Philosophies of Governance

The contrast with the United States is not merely a transatlantic curiosity. It reflects a genuine disagreement about how regulation should relate to technology it cannot fully anticipate.

The current US administration has been consistent in its position. Executive Order 14365, issued in December 2025 under the title Ensuring a National Policy Framework for Artificial Intelligence, frames the issue in direct terms: AI innovation must not be constrained by "onerous and excessive" rules, and US policy should sustain "global AI dominance" through "a minimally burdensome national framework." The America's AI Action Plan emphasizes permissive conditions for AI development and a streamlined federal position designed to prevent fragmented state-level regulation from creating compliance complexity. There is no binding federal AI statute. That is a deliberate posture, not a legislative gap.

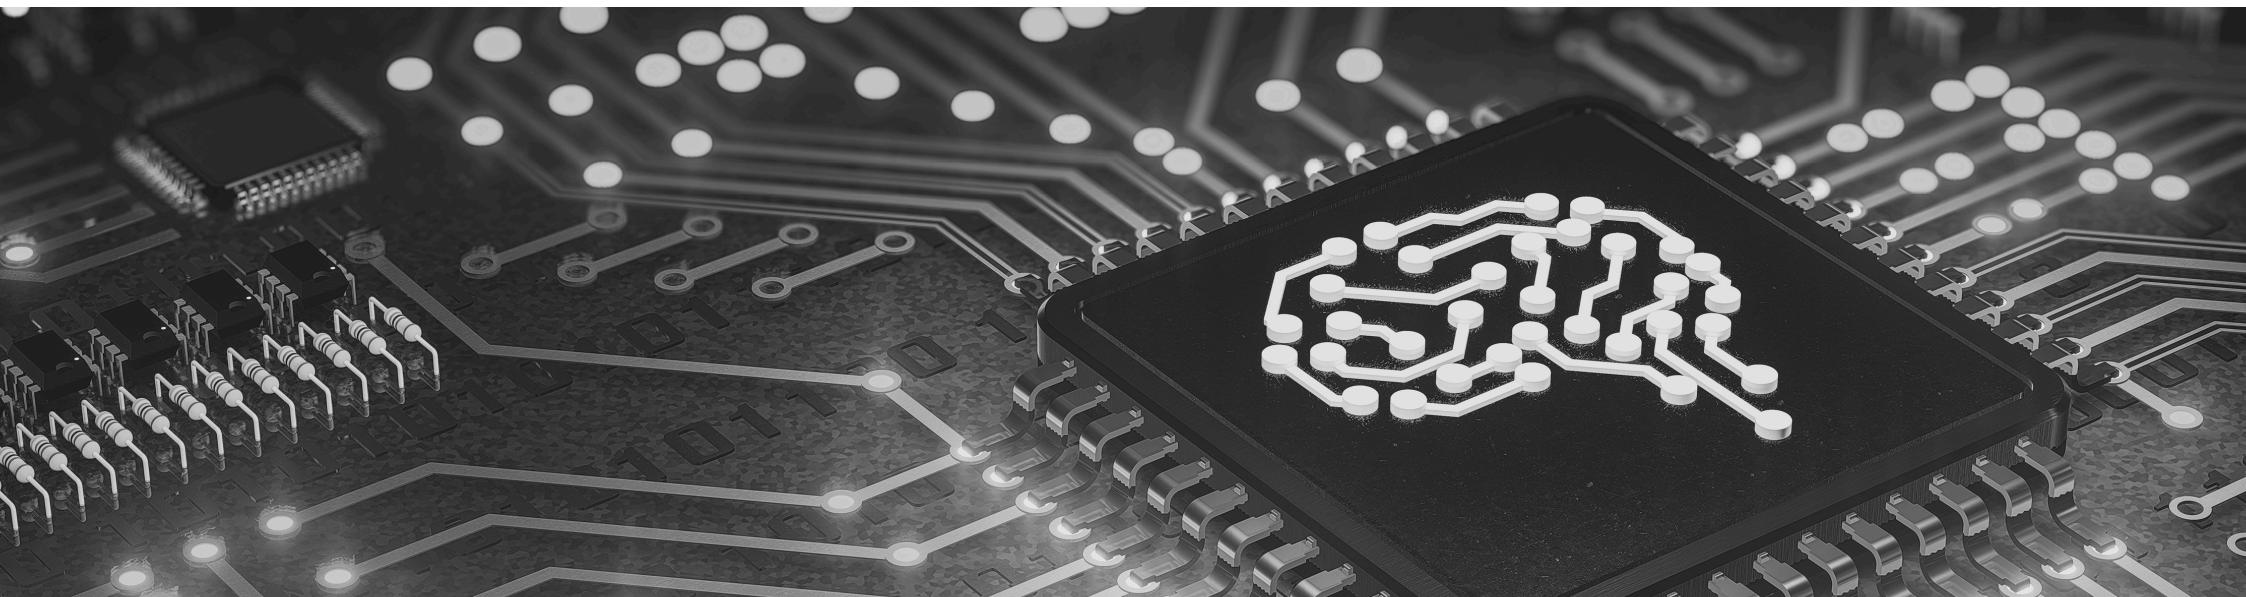

Beneath that political framing, the NIST AI Risk Management Framework operates from a different conceptual foundation. Where the EU AI Act asks what category a system falls into and what the checklist requires, the NIST RMF asks what risks a system poses in its actual operational context and what governance is proportionate to those risks. Its four functions – Govern, Map, Measure, Manage – are voluntary, scalable, and sector-agnostic. They do not presume to enumerate the uses of AI in advance.

This is not an argument without weaknesses. A voluntary framework produces voluntary compliance. In sectors where AI systems affect individuals who have no direct relationship with the deploying organization, voluntary governance and mandatory accountability are not interchangeable. The EU's insistence on enforceable rights and independent oversight reflects lessons learned from a decade of experience with GDPR: that transparency obligations without enforcement teeth tend to generate privacy notices rather than privacy protection.

Both observations can be true simultaneously. An industry survey of EU and UK tech SMEs found that six in ten report delayed access to frontier AI models, 58% say regulation has slowed product launches, and nearly half face measurably higher costs. Initial compliance costs for a single high-risk AI system under Annex I are estimated to be between 320,000 and €600,000 which are the figures that can erase 40% of a smaller firm's profits. The question for regulators is not which philosophy is correct in the abstract, but which combination of mandatory standards and adaptive governance produces systems that are genuinely safe, rather than systems that are compliantly documented.

The Case for a Context-Based Layer

The Act remains primarily categorical, but it is not purely binary. Article 6 allows a limited no-significant-risk derogation for some Annex III systems, subject to documentation and registration, while profiling systems remain automatically high-risk.

If it falls into the high-risk category, the full suite of high-risk obligations applies regardless of how the system is actually deployed, by whom, at what scale, or in what operational context. A large hospital deploying an AI diagnostic tool that processes thousands of patients daily and a small clinic piloting a scheduling assistant that touches clinical data face, in principle, the same compliance architecture.

The regulator's role is ex-post, triggered by complaints, market surveillance, or incident reports. The documentary layer (i.e. technical files, quality management systems, fundamental rights impact assessments) is procedurally thorough. Its actual protective value depends on whether the documents reflect operational reality.

A framework that supplements categorical risk classification with context sensitivity would ask different questions: not only is this system high-risk under Annex III, but how is it deployed, by whom, over how many people, with what human oversight in practice, and what feedback mechanisms exist when it errs. This is closer to how the GDPR's data protection impact assessment functions: context-dependent and proportionate rather than uniform. The EU AI Act already contains elements of this logic in its regulatory sandboxes and its recognition that different deployers have different capacities. What it lacks is a structural mechanism for translating context into proportionate obligation.

This gap is particularly relevant for jurisdictions now drafting their own frameworks. An AI governance architecture that assigns compliance duties according to the actual risk profile of a deployment, rather than the categorical label of its use case, would be more burdensome for genuinely risky applications and significantly less burdensome for the majority of deployments that carry limited systemic risk. It would also be more durable. A regulation amended before its primary obligations have applied is, at minimum, an argument for building differently next time.

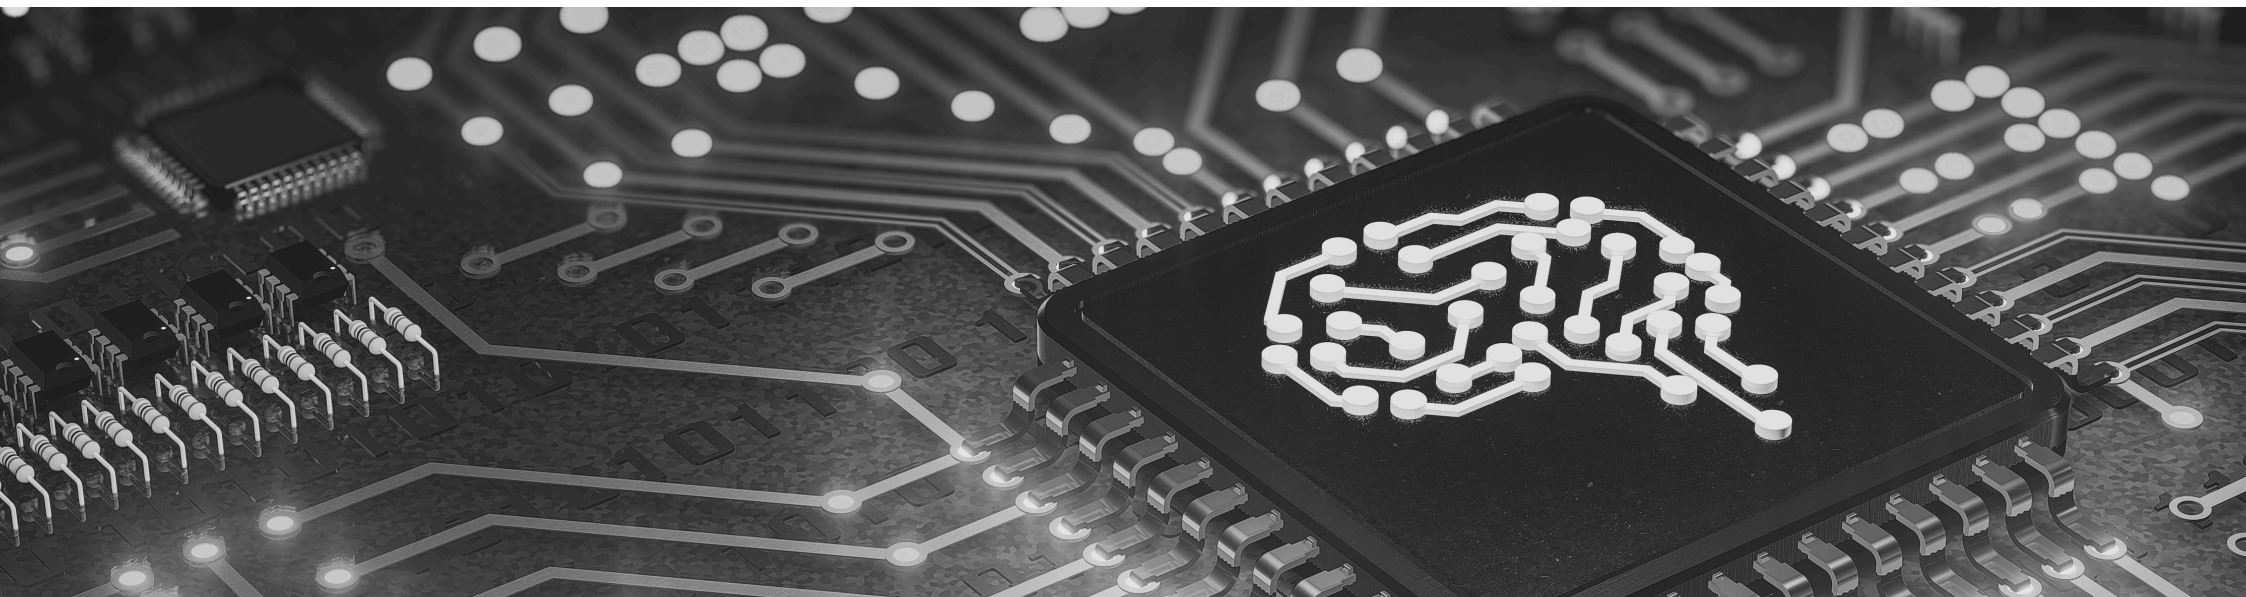

Practice for Deployers

For organizations currently deploying, or planning to deploy, AI systems within the EU market, today's deal changes the clock, but not the direction.

The AI literacy obligation under Article 4 already applies. The 2 August 2026 deadline remains relevant for Article 50 transparency obligations, while the provisional agreement sets 2 December 2026 as the deadline for the specific watermarking obligations for AI-generated content.

For Annex III high-risk systems, the December 2027 deadline should not be read as two additional years of inaction. Risk management systems, technical documentation, data governance, fundamental rights impact assessments where applicable, and human oversight procedures that function in practice rather than on paper take months to build properly. Organizations that begin that work now will have time to do it well.

Those extensions should not be treated as a pause in preparation. Risk management, technical documentation, data governance, human oversight, incident escalation and fundamental-rights impact assessment processes take time to build properly.

A Note on Adaptive Regulation

The revisions made to the EU AI Act in 2025 and again today are not a sign of failure. They are a sign that regulating a technology mid-deployment, without settled technical standards, in a market that is changing faster than the legislative cycle, requires continuous adjustment.

Geoffrey Hinton, whose foundational work on neural networks made much of modern AI possible, has observed that systems now entering enterprise deployment are acquiring capabilities that were not part of their original design – and that the tools for understanding, predicting, and constraining those capabilities are still being developed.

The governance challenge this poses is not primarily one of documentation. It is one of operational insight: understanding, in real time, what a deployed system is doing, why, and whether any person in the organization has the knowledge and authority to intervene if the answer becomes uncomfortable.

Some of the surveys suggest that around 40% of enterprise applications will integrate AI agents by the end of 2026. A 2026 industry survey found that 63% of organizations cannot enforce purpose limitations on their AI agents, 60% cannot terminate a misbehaving agent, and 55% cannot isolate AI systems from the broader network. A compliance document does not stop an autonomous agent from acting outside its intended scope.

Regulation that builds organizational capacity, encourages transparency about limitations, and creates proportionate incentives for genuine oversight rather than procedural compliance would be adaptive in the sense that matters. The EU AI Act points toward that goal. The distance between the pointing and the arriving is the most important gap in the current framework, and no extension of deadlines closes it.

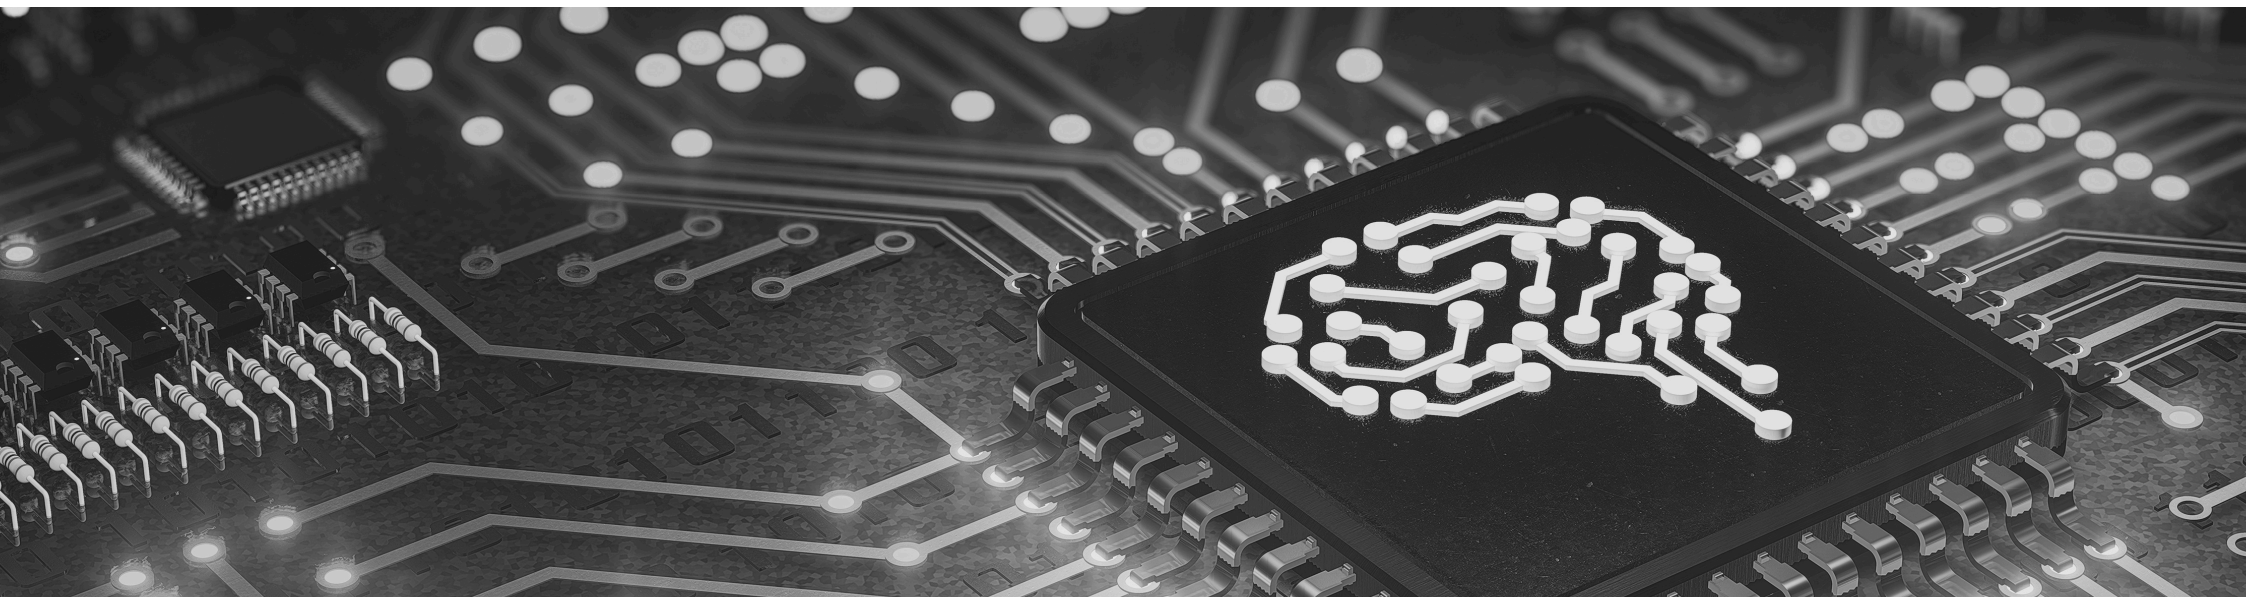

Closing Thought

The European Union built the world's first comprehensive AI regulation on a principle that is fundamentally sound: that technology affecting health, safety, and fundamental rights should be subject to independent oversight and democratic accountability before it causes harm rather than after. The principle was right. The implementation has been, to put it charitably, iterative.

The paradox is that Europe ends up with rules that are detailed enough to impose heavy compliance burdens, but not adaptive enough, or updated quickly enough, to keep pace with how technology develops. That combination can discourage innovation without improving safety in practice. Paper compliance and operational safety are not the same. The latter depends on the right framework and the right architecture, regardless of whether the deadline is August 2026 or December 2027.