

BEYOND RISK TIERS

Regulating What You Cannot Inventory — The AI Act Meets the Enterprise

The discovery programme hiding inside Europe's AI regulation

Giovanni Leonardi

September 2024



Regulating What You Cannot Inventory — The AI Act Meets the Enterprise

Portfolio · No. 105

Originally written September 2024 | Re-edited and re-rendered for the WhereBeyond Edition,
2025

*"The compliance artefact was complete; the
compliance reality was not."*

1. The Rhyme Is Structural

The compliance briefing follows a pattern I have now watched twice. A regulation arrives — sweeping, principled, risk-based — and the first slide deck frames it as a legal workstream. A policy team is assembled. Templates are drafted. Controls are mapped to articles. And then someone asks the foundational question that the entire programme depends on: what, exactly, are we governing? In the silence that follows, the real work begins.

The EU AI Act entered into force five weeks ago. Across the enterprises I observe, the response is falling into the same two camps that GDPR produced six years earlier: those treating the regulation as a legal exercise, and those who have recognised that the regulation is, in practice, a discovery exercise. The second group will comply. The first will produce policies for systems they cannot find.

The parallel with GDPR is not approximate — it is structural. GDPR required organisations to demonstrate lawful processing of personal data. The implicit prerequisite was knowing what personal data you held, where it lived, who processed it, and on what basis. This sounded elementary. It was not. The typical enterprise in 2017 held personal data in places its data protection officer had never mapped: legacy CRM exports on shared drives, customer service recordings in regional systems, analytics platforms ingesting behavioural data under contracts nobody in legal had reviewed. The regulation asked for governance. The real programme was cartography.

The AI Act repeats the pattern at higher stakes. Its risk-tiered obligations — prohibited practices, high-risk system requirements, transparency duties for general-purpose AI — all presume a prior answer to the same foundational question: which of our systems use AI, and what do they do?

2. Finding Things That Do Not Announce Themselves

What makes the AI inventory harder than the data inventory is not the technology — it is the visibility. Personal data, for all the difficulty of mapping it, at least had the decency to

be recognisable as data. It sat in databases and files and could, with effort, be traced. AI systems are frequently invisible to the people who depend on them.

Consider what a mid-sized financial services firm actually runs today. The vendor-supplied fraud detection system is obviously AI — it appears in procurement records and has an identifiable model owner. But the credit-scoring module embedded in the core banking platform? That uses machine learning, but it arrived as a feature update, not a procurement decision. The customer-service chatbot is visible; the sentiment-analysis layer feeding its routing logic is not. The marketing team's propensity model lives in a spreadsheet that a departed analyst built two years ago, and nobody is entirely certain how it works, only that it does. The HR department's CV-screening tool was adopted by a regional office under a team subscription and never went through central IT.

This is not negligence. It is the natural consequence of AI becoming infrastructure — embedded, distributed, and frequently acquired as a capability within something else rather than as a discrete system. The Act defines AI systems broadly enough to catch most of these, but the enterprise reality is that AI is a gradient, not a category. The boundary between a rule-based algorithm and a learned model, between a configured threshold and a trained classifier, is not always obvious to the people operating them. The Act's obligations attach regardless of whether someone in the organisation has thought to call the thing artificial intelligence.

3. The Self-Reporting Trap

The organisations treating the AI Act as a legal workstream are producing a familiar artefact: a compliance framework that is internally coherent and operationally hollow. Policies are drafted for high-risk AI systems. Risk assessments are templated. Governance boards are constituted. But the policies govern a register that was compiled by asking business units to self-report their AI usage — a method that reliably captures the systems people know they have and misses the ones they do not.

The register built on self-reporting describes perhaps sixty per cent of an organisation's actual AI landscape. The remaining forty per cent surfaces through incidents, audit findings, and questions nobody thought to ask. This is not a compliance gap — it is a governance fiction.

We watched this play out under GDPR. The organisations that built their compliance programme on self-reported data inventories spent eighteen months producing records-of-processing-activities that described a fraction of their actual data landscape. The remainder surfaced over the subsequent three years, usually through incidents, subject-access requests that could not be fulfilled, or audit findings that revealed processing nobody had declared. The compliance artefact was complete; the compliance reality was not.

The AI Act's enforcement timeline is more forgiving than GDPR's was — full application for most high-risk obligations does not arrive until August 2026 — but the inventory

problem is arguably less forgiving. You can discover a previously unknown personal-data processing activity and bring it under governance retrospectively with relatively contained consequences. Discovering a previously unknown high-risk AI system after it has produced a discriminatory hiring decision, a wrongful credit denial, or a flawed risk assessment is a different conversation entirely.

4. Cartography Before Policy

The firms that recognised GDPR as a discovery exercise rather than a documentation exercise took a different approach: they began with the data, not the policy. They mapped flows before they drafted privacy notices. They audited processing activities before they constituted governance boards. The policy followed the map, not the other way around.

The AI inventory programme follows the same logic, adapted for the particular difficulty of finding things that do not announce themselves as AI. In practice, it moves through three layers, each designed to catch what the others miss.

The first layer is procurement and vendor management — not because that is where all AI lives, but because it is where a structured record already exists. Every vendor contract and SaaS agreement signed in the last five years is a potential AI surface. The question to ask is not did we buy AI? but do any of our vendor products use machine learning, automated decision-making, or algorithmic processing? Framed that way, the question catches the embedded models that a narrower framing misses. Most vendor contracts from the last three years will answer it in their updated terms of service; for the rest, a direct inquiry to the vendor is warranted. The output is the first layer of the register: known, externally sourced AI.

The second layer is the technology estate — the platforms, tools, and infrastructure the organisation builds and operates. This is where shadow AI lives: the Python scripts that a data team wrote for demand forecasting, the RPA bots that make classification decisions their designers would not call AI, the Excel workbooks with regression models that have become operationally critical without ever being recognised as such. The search method here is technical, not administrative: scanning code repositories for model-loading libraries, reviewing automation platforms for decision logic, and — critically — asking the people who build things, not only the people who manage them. A developer will tell you about the model they trained; their line manager may not know it exists.

The third layer is the business process layer. Some AI systems are invisible not because they are technically obscure but because they were adopted as business decisions, not technology decisions — a team subscription to a scheduling tool that uses algorithmic optimisation, a recruitment platform with automated candidate screening, a customer-engagement tool that personalises content using behavioural models. These do not appear in the technology register because they were never registered as technology. Finding them requires walking the business process, not scanning the infrastructure.

The output of the three layers is not a spreadsheet. It is a portfolio view — a living register that maps each identified system to the Act's risk tiers, records its provenance, its data dependencies, its decision-making role, and its owner. The register is the foundation on which risk assessments, conformity procedures, transparency obligations, and

human-oversight arrangements can actually be built with confidence. Without it, those obligations are theatre.

5. What the Map Reveals

The firms that took the discovery approach to GDPR found something they had not expected: the data-mapping exercise revealed their operational reality with a clarity that no prior initiative had achieved. They discovered redundant systems, ungoverned data flows, and processing activities that contradicted their own architectural diagrams. The compliance programme, almost accidentally, became an operational improvement programme.

The AI inventory will do the same, and more so. The enterprise that catalogues every system in which a model makes or informs a decision will, for the first time, understand what its business actually runs on. It will discover which decisions are genuinely automated and which it only believes are human. It will find the single points of failure — the departed analyst's spreadsheet model, the vendor algorithm that nobody in the organisation can explain, the classification system that has drifted from its training data without anyone monitoring it — and it will be able to address them before they produce the harm event that the regulation exists to prevent.

The AI Act is a regulatory obligation. But the inventory it demands is something more valuable: the first honest look most enterprises will take at the machinery that now drives their operations.

The organisations that understand this will not merely comply. They will, incidentally, learn what their business actually does. We have seen this rhyme before — the question is whether we will recognise the verse this time, or wait, again, for the chorus to teach us what the opening already said.