THE CLOUD PEOPLE
An NTT DATA Business Solutions Company

AN EXECUTIVE BRIEF

The AI you've bought may be illegal to deploy.

How European enterprises can ensure their AI investments are **effective**, **efficient**, and **compliant by design** — before the regulatory calendar runs out.

BY

The Cloud People

Autonomous IT CoE

ServiceNow Practice · Founded 2019

MAY 2026

REF · BRIEF-AI-01

01 ✦ THE SITUATION

More time, less margin.

On 7 May 2026, EU negotiators reached a provisional agreement — the Digital Omnibus on AI — to postpone the core enforcement deadline for high-risk AI systems from August 2026 to December 2027. It was widely reported as relief. In practice, it is not.

Conformity assessments for high-risk AI take six to twelve months. Harmonised standards and guidance documents are not yet final. Most European enterprises have neither the documentation, the data governance, nor the service-level visibility required to demonstrate compliance. A 16-month extension to December 2027 is, for most organisations, exactly the time needed to do the work properly — assuming work begins now.

Meanwhile, several pieces of the regulatory picture remain unchanged and immediately relevant to the board:

- **EU AI Act prohibited practices and AI literacy obligations** have been enforceable since February 2025. The **€35 million / 7% global turnover** penalty framework is already live.
- **DORA** has been in force across European financial services since January 2025, with administrative fines of up to **2% of total annual worldwide turnover** (and higher in some member states). Roughly half of in-scope institutions are estimated not yet fully compliant.
- **EU AI Act transparency obligations** (Article 50) for AI-generated content apply from December 2026 — seven months from now.
- **High-risk AI obligations** (Articles 9–15) apply from December 2027 for use-based systems and August 2028 for product-embedded systems. The substance of those obligations has not been relaxed.

> **The delay is not relief. It is a tighter calendar disguised as more time. The organisations that read it as breathing room will discover, in late 2027, that they had exactly the runway they needed — and they have already burned half of it.**
>
> **The Cloud People · Autonomous IT CoE**

02 ✦ THE CONSTRAINT

What the regulator will actually ask.

The major enterprise platforms have restructured to bundle AI as a baseline rather than an option. Customers renewing in 2026 are increasingly paying for autonomous AI whether they activate it now or not — so the operative question is no longer whether to deploy AI, but whether the organisation can prove the AI it deploys is governed and accountable. For most European enterprises today, the honest answer is no.

Strip the EU AI Act and DORA down to their operational essence, and they converge on the same set of questions. Both regulations require organisations to demonstrate, on demand and with evidence, four things about every AI system or critical ICT service in scope:

- **What does this system do**, and which business function does it serve?
- **Who is accountable** for it — by name, not by team or queue?
- **What is connected to it**, and what depends on it failing or behaving incorrectly?
- **What state is it in** — in production, under review, decommissioned, modified since its last assessment?

These are not novel questions. Every IT organisation has been asked some version of them for years. What is new is that the answers must now be evidenced, audit-ready, and consistent across the enterprise — and that AI systems, which act at machine speed and in parallel, expose any inconsistency immediately.

The honest diagnosis from our work with European enterprises is that most organisations cannot produce these answers cleanly today. Ownership fields are populated with generic queues. Service relationships are partial or stale. Lifecycle status is approximate. Data classifications are missing or inconsistent. None of this prevented the business from functioning when humans were in the loop, interpreting context and filling gaps. With AI systems acting autonomously and regulators demanding documentation, both safeguards disappear at once.

The structural point follows directly: **AI compliance is not an AI problem. It is a service-model problem that AI exposes ruthlessly.** Each of these questions has a precise answer in a well-modelled service environment, and none of them has a precise answer where the configuration database is a list of servers and the service model is an afterthought. The legal and audit obligations the EU is imposing in 2026–2028 will be met or missed not by AI procurement decisions, but by the maturity of the operational data substrate enterprises have been quietly underinvesting in for a decade.

03 ✦ THE DECISION

Three questions for the next board meeting.

The argument of this brief reduces to three questions a CEO or CIO should be able to answer — or should be uncomfortable not being able to answer — before the next renewal cycle.

1. Can we name the owner?

For every business-critical application, every ICT service supporting a critical function, and every AI system in scope for EU AI Act high-risk classification — can we name a single accountable person, not a team or a queue? If not, EU AI Act Article 14 (Human Oversight) and DORA Article 28 third-party register submissions cannot be defended.

2. Can we map the impact?

If an AI system or a critical ICT service fails or behaves incorrectly, can we identify — within minutes, not days — which business processes, customers, and revenue streams are affected? If not, neither risk management (Article 9) nor DORA's operational resilience requirements can be evidenced.

3. Can we reconstruct the state?

If a regulator asks us to produce technical documentation for an AI system in production — what version is running, what data it uses, what has changed since its last assessment — can we do so in days rather than months? If not, Article 11 (Technical Documentation) and Article 12 (Record-Keeping) will not survive scrutiny.

Where the honest answer to any of these is no, the gap is not in AI capability — it is in the service model that AI and compliance both consume. Closing that gap is unglamorous work: ownership cleanup, relationship mapping, data classification, and lifecycle hygiene. It does not have to be finished before any AI is switched on — but on the systems where the answers are no, it is the constraint that decides whether AI there can be deployed safely and defended to a – regulator.

Where to start

The mistake to avoid is treating this as a sequence — foundation first, AI later. Most of the AI capabilities an enterprise is already paying for deliver value on low-risk, well-understood

processes where the service model is good enough today. Those should be switched on now. The work is to know which systems are safe to activate immediately and which sit behind a real compliance or quality gap — and then to run activation and hardening in parallel, not one after the other. The Autonomous IT Centre of Excellence at The Cloud People organises this around three workstreams that overlap by design:

1. **Readiness Assessment.** A short, structured diagnostic that maps your AI ambitions and your current service model against the specific articles of the EU AI Act and DORA that apply. The output is a triage: what is safe to activate today, what is exposed and needs hardening first, and what the realistic timeline looks like for each.
2. **Activate where it's safe.** Switch on the AI capabilities that deliver value now — on the processes where the service model already supports them. Value starts accruing immediately against the spend you have already committed, rather than waiting on a multi-month foundation programme.
3. **Harden in parallel.** Close the ownership, relationship, classification, and lifecycle gaps on the higher-risk and higher-exposure systems — producing the evidence the regulator and the audit committee require — while the safe-to-activate capabilities are already in production.

The discipline is not in delaying AI. It is in matching the pace of activation to the maturity of the underlying data, system by system — moving fast where the foundation supports it, and hardening deliberately where it does not. Done this way, value and compliance advance together, and neither waits on the other.

€35M

Maximum EU AI Act penalty, or 7% of global turnover (prohibited practices live since Feb 2025)

2%

DORA administrative fine ceiling, of worldwide turnover, enforceable since January 2025

~18m

Realistic window for conformity readiness before the December 2027 deadline

About this brief. This is a short executive document. The full reasoning, regulatory references, and platform-specific analysis are available in our accompanying whitepaper, "*You can't govern what you can't see*" (CSDM-AI-02, May 2026). Regulatory references draw on the EU AI Act final text (Articles 4, 5, 9–15, 50), the Digital Omnibus on AI provisional agreement of 7 May 2026, Regulation (EU) 2022/2554 (DORA), and the published guidance of the European Commission and the European Supervisory Authorities as of May 2026. Final adoption of the Omnibus is expected before 2 August 2026.

18 months, used well.

This perspective comes from our **Autonomous IT CoE** at The Cloud People — where we help European enterprises switch on the value of AI on ServiceNow while bringing CSDM maturity, EU AI Act, and DORA obligations forward in parallel.

> ↗ **NEXT STEP**
>
> A CSDM Readiness Assessment, mapped against the specific articles of the EU AI Act and DORA that apply to your enterprise. We tell you what is defensible today, what is exposed, and what the realistic timeline is — in days, not months.

THE CLOUD PEOPLE
An NTT DATA Business Solutions Company

thecloudpeople.com

Autonomous IT Centre of Excellence
ServiceNow Practice · Focus · Fair · Future · Fun