REGULATORY INSIGHT

Procuring the Black Box: AI Vendors and the Limits of Third-Party Due Diligence

Financial entities are, increasingly, buyers of artificial intelligence rather than builders of it. Credit decisioning, fraud analytics, document review, customer service, and underwriting support are routinely delivered by third-party AI systems — often built on foundation models the firm neither trained nor controls. This is a rational commercial choice. It is also a distinctive legal problem, because an AI vendor is opaque in ways a traditional ICT supplier is not.

When a firm procures a conventional software service, due diligence has a defined object: the service can be specified, tested, audited, and, in principle, replaced. When a firm procures an AI system, the object of due diligence is partly unknowable. The model's behaviour emerges from training data the firm cannot see, weights it cannot inspect, and update cycles it does not control. Yet under EU law, the firm carries the regulatory consequences of that system as if it understood it fully. This article examines the seam where AI procurement meets operational resilience and AI governance — and where the legal function now has to do its hardest work.

Two Regimes, One Vendor

An AI procurement in a regulated financial entity is governed by two frameworks at once.

The Digital Operational Resilience Act treats the AI vendor as an ICT third-party service provider. That brings the full third-party apparatus into play: the arrangement must appear in the Register of Information; the contract must satisfy DORA's minimum contractual requirements, including access and audit rights and provisions on subcontracting and exit; and, where the AI supports a critical or important function, the firm must assess concentration risk and substitutability.[1]

The AI Act, separately, treats the firm as a deployer of an AI system — and, where the use case is high-risk, imposes deployer obligations under Article 26: using the system in line with the provider's instructions, ensuring meaningful human oversight, monitoring operation, and retaining logs.[2] For financial institutions, some of these obligations are designed to dovetail with existing financial-services governance rules, and several relevant deadlines have provisionally been extended through the Digital Omnibus on AI political agreement of 7 May 2026, which remains subject to formal adoption.[3] But the obligations do not disappear; they accumulate on top of the DORA requirements.

The two regimes land on the same contract, the same vendor, and the same procurement decision. They are not alternatives. They are layers.

The Documentation Gap

The central practical difficulty is documentation. Both regimes assume the firm can see enough of the system to govern it. Under the AI Act, the provider must prepare detailed technical documentation in line with Article 11 and Annex IV, and must make information about the system, including instructions for use, available to deployers under Article 13 — the inputs deployers need to oversee a high-risk system and interpret its output.[4] DORA, separately, presumes audit and access rights meaningful enough to give assurance over a service supporting critical functions.

Foundation-model vendors, as a matter of commercial reality, resist exactly this. Training data composition, model architecture, evaluation results, and known failure modes are treated as proprietary. Standard AI procurement contracts frequently concede none of the documentation, audit access, or model-change transparency that the two regimes contemplate — a gap that, as supervisors have begun to observe across ICT contracts generally, remains widespread and unremediated.[5]

This is a legal-function problem, not a procurement-function problem. It cannot be solved by accepting the vendor's paper. It has to be negotiated: the contract must require technical documentation sufficient for the firm's own compliance obligations, audit or assured-assurance rights, notification of material model changes, and cooperation in incident handling. Where a vendor will not move, that refusal is itself a finding — one the firm should record, escalate, and weigh in the decision to proceed.

Who Is the Provider, and Who Is the Deployer?

A subtler risk concerns role allocation. The AI Act distinguishes between the provider of a system and its deployer, and assigns markedly heavier obligations to providers. A firm that merely uses a vendor's system as supplied is a deployer. But a firm that fine-tunes a model, repurposes it for a use the vendor did not intend, or substantially modifies it can find itself recharacterised as a provider of a high-risk system under Article 25 — read with the Article 3(23) definition of “substantial modification” — inheriting conformity assessment and documentation obligations it never planned for.[6]

This is not a remote scenario. The line between configuring a system and modifying it is genuinely fine, and it is crossed in ordinary engineering work. The legal function should understand, before contract signature, exactly what the firm intends to do with the model, and should price the regulatory consequences of crossing that line into the build-versus-buy decision.

When the AI Vendor Is Also a Critical Third Party

The concentration problem compounds the procurement problem. The largest providers of enterprise AI are the same hyperscale cloud and technology firms that now appear on the first list of designated critical ICT third-party providers under DORA.[7] A firm procuring AI capability from one of these providers may

be deepening a dependency on an entity already under direct EU oversight — and already a node of systemic concentration.

That does not make the procurement wrong. But it means the firm should treat the AI arrangement as part of its concentration analysis, not as a standalone contract. The relevant question is not only “is this AI service resilient?” but “what does this arrangement add to our exposure to a provider we already cannot easily replace?”

Change You Cannot See

Traditional ICT change management assumes change is discrete and visible: a version is released, the firm tests it, the firm adopts it. AI systems do not always behave that way. Models are retrained and updated, sometimes silently, and behaviour can shift without any change the firm initiated or observed.

Both DORA's change management discipline and the AI Act's oversight obligations assume the firm knows what it is running. A model that changes underneath a deployer undermines that assumption. Contracts therefore need to do unfamiliar work: requiring advance notification of material model changes, defining the firm's right to evaluate or reject an update before it reaches production, and preserving the firm's ability to demonstrate, at any point, which version produced a given decision.

Incidents That Wear Two Hats

An AI failure can be a regulated event under both frameworks simultaneously. A malfunction that disrupts a critical function may be a major ICT-related incident reportable under DORA. The same malfunction, if the system is high-risk, may be a serious incident reportable under Article 73 of the AI Act — to a different authority, on a different tiered timeline (2 days for widespread infringement or critical-infrastructure disruption, 10 days where the incident has caused death, and 15 days for other serious incidents), against a different template.[8] Where a financial-supervision authority has been designated as the market surveillance authority for AI under Article 74(6)–(7), the same regulator may receive both notifications.

A firm's incident response framework should anticipate this dual character before an incident occurs. The classification logic, the escalation paths, and the reporting playbook all need to account for the possibility that one event triggers two regimes — and the AI vendor's contractual obligation to cooperate in both should be secured at the outset, not negotiated mid-incident.

Exit From a Black Box

Finally, exit. DORA expects firms to maintain credible, tested exit strategies for arrangements supporting critical functions.[9] Exit from an AI vendor is conceptually harder than exit from a conventional service. The model, its weights, the prompts and embeddings tuned to it, and the institutional knowledge built around

its behaviour are not readily portable. A firm cannot simply lift a decisioning capability from one foundation model and set it down on another and expect equivalent performance.

A realistic exit strategy for an AI arrangement is therefore less about portability and more about substitutability: maintaining the firm's ability to fall back to an alternative model, a simpler rules-based process, or a manual process, and testing that fallback rather than assuming it. An exit plan that exists only on paper is, here as elsewhere, the kind of gap supervisors are now looking for.

Practical Priorities for Financial Entities

- **Run AI procurement through both regimes from the start.** Map every prospective AI arrangement against DORA's third-party requirements and the AI Act's deployer obligations before contract terms are agreed.
- **Negotiate the documentation gap; do not accept the vendor's paper.** Require technical documentation sufficient for your own compliance, audit or assurance rights, and model-change notification. Record vendor refusals as findings.
- **Decide your provider-or-deployer status deliberately.** Understand whether fine-tuning or repurposing will recharacterise the firm as a provider under Article 25, and account for that before the build-versus-buy decision is made.
- **Fold AI arrangements into concentration analysis.** Where the AI vendor is also a critical third party, assess the combined exposure, not the contract alone.
- **Build incident playbooks for dual-regime events.** Ensure classification and escalation logic handles incidents reportable under both DORA and the AI Act, and tracks the tiered AI Act timelines.
- **Test, do not assume, your AI exit.** Treat exit as substitutability — a tested fallback to an alternative or simpler process — rather than portability.

Conclusion

Buying AI is not the same as buying software, and the legal function is where that difference is felt most sharply. Two regulatory frameworks now converge on every significant AI procurement a financial entity makes, and both assume a level of visibility into the vendor's system that the vendor is commercially reluctant to provide. The firms that manage this well will not be the ones that procure AI fastest, but the ones that treat the contract as the principal control point — closing the documentation gap, fixing role allocation, and securing change and incident cooperation before signature. The black box cannot be made fully transparent. It can, with disciplined drafting, be made governable.

References

1. Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 (DORA), in particular Chapter V on the management of ICT third-party risk; DORA Article 3(21) (definition of ICT services).
2. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (the EU AI Act), Article 26 (obligations of deployers of high-risk AI systems).
3. Council of the European Union, press release, *"Artificial intelligence: Council and Parliament agree to simplify and streamline rules"* (7 May 2026); the agreement remains subject to formal adoption. See also Travers Smith, *"EU Agrees to Delay Key AI Act Compliance Deadlines"*.
4. EU AI Act (n 2), Article 11 and Annex IV (technical documentation prepared by the provider) and Article 13 (transparency and provision of information to deployers of high-risk AI systems).
5. EU AI Act (n 2), Article 13 documentation requirements; BaFin/Deutsche Bundesbank, *IT-Aufsicht im Finanzsektor 2025 — Vortrag 1* (December 2025), on the widespread non-remediation of ICT contracts to DORA's minimum contractual requirements.
6. EU AI Act (n 2), Article 25 (responsibilities along the AI value chain), read with Article 3(23) (definition of "substantial modification"); see also the European Commission's draft guidance on the provider/deployer line (April 2025).
7. EBA, EIOPA and ESMA, joint press release, first list of designated critical ICT third-party service providers under DORA (18 November 2025).
8. EU AI Act (n 2), Article 73 (reporting of serious incidents — tiered timelines of 2, 10 and 15 days) and Article 74(6)–(7) (designation of financial-sector authorities as market surveillance authorities for AI in connection with regulated financial services); see also K&L Gates, *"EU and Luxembourg Update on the European Harmonised Rules on Artificial Intelligence — Recent Developments"* (20 January 2026).
9. DORA (n 1), provisions on exit strategies for ICT services supporting critical or important functions; BaFin/Deutsche Bundesbank (n 5), on supervisory attention to exit planning.
