FORENSIO · WHITE PAPER · JULY 2026

The AMLA-Era Decisioning Stack

What regulators now expect from AML decisioning, where the market stands in 2026 — and the questions to ask any vendor before you sign.

1 · The regulatory clock is no longer abstract

Four regimes converge on AML decisioning within twenty-four months, and each one asks for something specific and testable:

Date	Instrument	What it demands of a decisioning system
Jan 2025 (in force)	EU DORA	Art. 30 contract content for every ICT provider, a register-of-information entry, audit and exit rights. Cloud gained no free pass — the chain of subcontractors is now the buyer's problem.
2 Aug 2026	EU AI Act — high-risk obligations	Art. 12: automatic event logging. Art. 14: effective human oversight by competent people — oversight that can be *demonstrated*, not asserted. Penalties reach 7% of global turnover.
~2027	EU AMLR single rulebook	Harmonised CDD and reporting; Art. 75 gives information-sharing partnerships a legal basis — with pseudonymisation mandated. Art. 77: delete personal data five years after the relationship ends.
2028	AMLA direct supervision	A single EU supervisor with inspection powers — and an appetite for artifacts, not narratives.
standing	US: SR 11-7 · NYDFS Part 504 · FinCEN	Model documentation and validation evidence; an officer **personally certifies** the monitoring programme every April 15; SAR e-filing runs on a published XML schema.
standing	DE: KWG § 25h · BaFin AuA	Data-processing monitoring systems are a statutory duty for credit institutions; supervision is documentation-first.

2 · Where the market actually stands

The leading platforms have converged on the same vocabulary — explainable AI, policy-bound agents, complete audit trails, regulator-ready outcomes. Three capabilities became genuine table stakes in 2025–26:

- **No-code rule management with pre-deployment testing** — compliance teams change detection logic without engineering tickets and test it against historical data first.
- **AI-assisted investigation** — copilots that summarise cases, gather adverse media and draft narratives.
- **Explanations attached to every alert decision** — plain-language rationale per case.

Yet underneath the shared vocabulary, three questions still separate the field — and they are precisely the questions the new regimes ask:

Question	Typical answer today	What the regulation actually rewards
Can you **re-run** last quarter's decision and get the same answer, byte for byte?	No — scores drift with models and data snapshots	Art. 12 logging + SR 11-7 validation both presume decisions that can be reconstructed exactly
When policy changes, do you know **exactly which past outcomes** would change?	A backtest on a sample — statistics with confidence intervals	Change governance that names the affected cases, clause by clause
Where does the **audit binder** come from?	Consulting engagements, per cycle	Documentation that regenerates from the system itself and can be hash-verified

The pattern behind all three: **probabilistic scoring engines cannot answer them, deterministic decisioning can.** A deterministic engine — same alert, same versioned policy, same decision, provably — turns reproducibility, exact replay and generated documentation from projects into properties.

3 · Human oversight that survives scrutiny

The known failure mode of "human in the loop" is rubber-stamping — a reviewer who clicks approve adds no oversight, and both GDPR Art. 22 case law and the AI Act's Art. 14 take aim at it. The emerging answer is **counterfactual context**: show the reviewer not just why the system concluded what it did, but *how close the call was* — which single fact or parameter would flip it. Exact counterfactuals require full re-execution under varied inputs; approximations (feature attributions) are the best a scoring engine can offer. This is a structural argument for determinism, not a feature comparison.

4 · Sovereignty stopped being a preference

DORA made every layer of the ICT chain contractual; banking secrecy regimes and GDPR transfer analysis price in every border the data crosses; the US CLOUD Act shadows any US-cloud dependency. An on-prem deployment answers the entire question class at once: locations — your premises; subcontractors — none; data return on exit — it never left. The evaluation question is no longer "cloud or on-prem?" but "**what does your vendor's answer to DORA Art. 30 look like, line by line?**"

5 · The evaluation checklist

Ask any AML decisioning vendor — including us — to **demonstrate live**, not describe:

- Re-run a historical decision. Is the result byte-identical, and how is that proven?

- Change a policy threshold. Show **exactly which** historical cases change outcome, and which merely change their cited justification.
- Show one decision's counterfactuals: what single variation flips it?
- Generate the model-governance / AI-Act documentation pack. How long did it take, and can its integrity be verified afterwards?
- Export a filing (goAML / FinCEN SAR). Is every sentence of the narrative traceable to evidence?
- Tamper with one byte of the audit trail. Does the system detect it, and how precisely?
- Show the DORA Art. 30 annex and the register-of-information sheet, pre-filled.
- Who can see a case's SAR status? Show the compartment, not the policy document.

Any vendor who can do all eight live deserves your shortlist. Most cannot; the vocabulary is converging faster than the engineering.
