Forensio FCC DECISIONING

WHITE PAPER

Anti-Black-Box: why a deterministic engine beats a risk score

How policy-bound decisioning gives model-risk and audit teams a trail they can actually validate — without slowing the analyst down.

Forensio · 2026 · On-prem explainable financial-crime decisioning

Financial-crime compliance is, at its core, a decision-making function. Yet most AI brought into the discipline outputs a *score* — "87% risk" — and leaves the actual decision, and its justification, to the analyst and a spreadsheet. For a regulated institution, that is the wrong artifact.

1. The problem with a risk score

A probability is not a decision, and it is not defensible on its own. When an inspector asks "why was this case closed?", three things must be reconstructable: the **evidence**, the **rule or policy** that was applied, and the **person** who approved it. A score answers none of them. Worse, an opaque model that produces the score cannot be meaningfully validated by a model-risk team, which turns an efficiency tool into an audit liability.

2. A different shape of decision

Forensio inverts the usual arrangement. AI agents investigate — they gather evidence, resolve entities, compare transactions to profile, and draft a narrative — but they do not decide. A **deterministic reasoning engine** takes the classified evidence and applies the bank's own rules, policy, typologies and risk appetite to certify what disposition is permissible. A human then approves under maker-checker.

> **The principle:** AI proposes; the deterministic engine permits; a human decides. Every conclusion is linked to a specific piece of evidence and a specific policy reference, and the whole chain is versioned.

3. Why model-risk and audit teams prefer it

CONCERN	BLACK-BOX SCORE	FORENSIO DETERMINISTIC ENGINE
Validation	Opaque; hard to validate (SR 11-7 / EBA)	Validate rules & policy you can read
Explainability	Post-hoc, approximate	Built-in: each clause is a citation
Reproducibility	Drifts with model	Same evidence + version ⇒ same decision

Audit trail	Reconstructed manually	Sealed, append-only decision record

4. The complete decision record

Every adjudication produces a sealed record containing the evidence, the rule and policy versions, the model version, the user, the timestamp, the rationale and the approval chain. It is append-only and exportable to the institution's case manager — the artifact an auditor actually wants, generated as a by-product of the work rather than reconstructed after the fact.

5. Conclusion

Explainability is not a feature to bolt on after a model ships; it is the governing constraint. By binding AI investigation to a deterministic, policy-linked engine and a human approval, Forensio turns the hardest objection to AI in compliance — "we can't defend it" — into the product's strongest property.

> **See it live.** The interactive demo and Investigator Console at **forensio.com** show a real alert resolved into a sealed, regulator-ready decision on synthetic data.