Forensio FCC DECISIONING · ON-PREM

SOLUTION BLUEPRINT · PUBLIC EDITION

On-prem, policy-bound decisioning for financial crime

How Forensio turns screening and transaction-monitoring alerts into defensible, regulator-ready decisions — inside the institution's own environment.

2026 · Forensio · forensio.com · hello@forensio.com

AI proposes. A deterministic, policy-bound engine permits. A human decides. Every decision is evidence-backed, reproducible and sealed — and it never leaves the institution's environment.

1 · What Forensio is

Forensio is an **on-premise, policy-bound AI decisioning platform for financial-crime compliance**. It sits *above* the screening and transaction-monitoring systems an institution already runs, ingests their alerts, resolves the entities and money onto **one evidence graph**, lets a team of **specialist AI agents investigate**, and then a **deterministic engine certifies the permissible disposition** under the bank's own policy. A human approves under maker-checker, and the whole chain is written to a **sealed, append-only decision record**.

The wedge: we don't replace your detection. We turn its alerts into defensible decisions — without a black box and without your data leaving the building.

2 · Where today's tooling falls short

Modern financial-crime platforms have converged on a common pattern language: unified case handling, entity resolution with network graphs, AI-assisted adjudication, real-time hybrid scoring, self-serve rules with sandboxes. Forensio builds on all of it — and closes four gaps that remain:

THE GAP IN THE MARKET	FORENSIO'S ANSWER
Cloud-first delivery; true on-prem / air-gap with zero egress is rare or retrofitted.	**On-prem is the backbone, not an option — data and models stay inside the institution.**
"Explainable AI" usually means a score plus a post-hoc explanation — hard for model risk to validate.	**A deterministic, policy-linked decision: reproducible, validatable (SR 11-7 / EBA), not approximated.**
Screening and transaction monitoring remain separate engines and queues.	**One evidence graph reasons across identity-based and behaviour-based risk in a single pass.**
Nobody audits the analyst's own decision before it closes.	**The Decision QA agent checks completeness, consistency and audit-readiness — and blocks indefensible closes.**

3 · The decisioning model

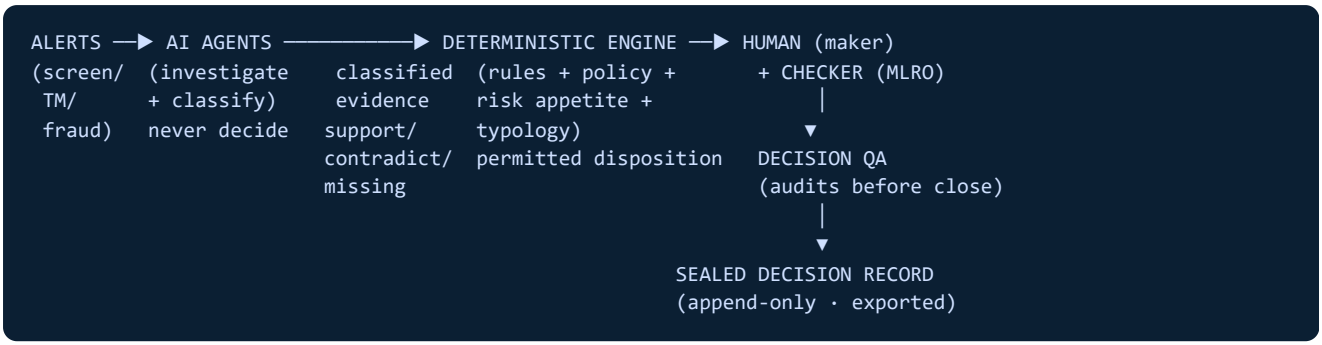

Why this beats a risk score: your model-risk team validates rules it can read; every conclusion cites a specific piece of evidence and a versioned policy clause; the same evidence and versions always produce the same decision; and the sealed record is generated as a by-product of the work — not reconstructed for the inspector.

4 · Built on proven, open frameworks

```
PRESENTATION   Investigator Console · Operations Dashboard · Decision Record · Network map
DECISIONING    Deterministic engine: DMN rules (GoRules/Drools) + BPMN (Camunda) · Decision QA
AI AGENTS      Locally served models (vLLM / Ollama) — investigate & classify, never decide
GRAPH & ER     Entity resolution (Senzing / Zingg) · property graph (Neo4j / Memgraph) · FollowTheMoney
DATA           Sanctions/PEP (OpenSanctions, mirrored locally) · customer/KYC · transactions · policy repo
PLATFORM       Kubernetes / OpenShift · on-prem / private cloud / air-gap · SSO/RBAC · full audit log
```

Every component is institution-grade, open and self-hostable — no proprietary monolith, no vendor lock-in. The stack is yours, inside your walls.

5 · One platform, every seat at the desk

SEAT	WHAT THEY SEE
Investigator	A three-pane console: risk-scored queue · case workspace (Customer 360, EDD checklist, interactive network, transactions, timeline) · AI copilot with per-agent findings.
Operations lead	A configurable dashboard: AI insights over the alert population, an auto-generated shift brief, alert-load heatmap, SLA and throughput widgets.
MLRO / Audit	The sealed decision record: evidence with stances, agent findings, policy citations, the engine's "permitted" certification, maker-checker approvals, hash seal.
Head of Compliance	Board view: automation rate, analyst-hours saved, cost avoided, SLA compliance and SAR output — with trends.

All four are live today as an interactive demo at **forensio.com** — on synthetic data, in the browser.

6 · From alert to sealed decision

Intake (alert lands, entity-resolved onto the graph) → **Triage** (risk-scored, tiered routing) → **Investigate** (agents classify evidence; structured checklist) → **Decide** (engine returns the permitted disposition with binding clauses) → **QA + approve** (Decision QA audits; checker approves) → **Record + export** (sealed record; SAR/STR, goAML) → **Learn** (outcomes feed dashboards and model-risk validation).

7 · Security & compliance posture

- **Zero data egress** — on-prem / private cloud / air-gapped; models served locally.
- **RBAC + SSO**, segregation of duties, enforced maker-checker.
- **Append-only, hash-chained records**; full audit log; versioned rules, policy and models.
- **Aligned to:** FATF risk-based approach · EU AMLA/AMLR · SR 11-7 · EBA guidelines · GDPR · goAML export. **On the roadmap:** SOC 2 Type II, ISO 27001.

8 · Start where the risk is lowest

> **Decision-Quality & Explainability Pilot (€75–150k):** a contained replay of 500–5,000 of your historical screening + TM cases, inside your environment, read-only. You receive a decision-quality and gap analysis, a

false-positive automation estimate and an analyst-time benchmark — with zero production risk.
hello@forensio.com