

▣ Forensio — Compliance Pack (EU variant)

Generated from live artifacts — versions, seals and evidence below are read from the running system, not written by hand. Verify this pack against its manifest hash.

1 · System card

- **System:** Forensio Core — on-prem, policy-bound FCC decisioning
- **Package version:** 0.1.0 · **engine:** 0.1.0 · **source commit:** 5093c99
- **Active policy:** v19 (checked-in baseline)
- **Decision path:** evidence graph → deterministic agents (investigate & classify) → deterministic engine (certifies the permissible disposition; fails closed) → maker-checker approval → Decision QA (fails closed) → sealed, hash-chained record
- **Core properties:** byte-identical reproducibility (same alert + versions ⇒ same seal); zero data egress (no outbound calls; optional webhook targets bank-internal endpoints only); policy changes only through replayed, sealed Studio activations

2 · EU AI Act mapping (Regulation (EU) 2024/1689)

Obligation	Mechanism in this system
Art. 9 — risk management	Policy Studio: every policy draft is replayed against the full sealed history; the exact per-case diff is reviewed and sealed before activation (no-replay-no-activation)
Art. 11 — technical documentation	this pack (generated per release) + docs/ARCHITECTURE.md + published API contract (<code>GET /schemas</code> , <code>/openapi.json</code>)
Art. 12 — record-keeping / logging	append-only, SHA-256 hash-chained decision ledger; every decision, policy activation and replay report is sealed; <code>GET /ledger/verify</code> re-derives the chain on demand
Art. 14 — human oversight	maker-checker on every disposition; the engine only <i>*permits*</i> — a human decides; Decision X-Ray shows the checker how close each call was (exact counterfactuals), making oversight demonstrably meaningful
Art. 15 — accuracy & robustness	deterministic engine proven by CI reproducibility tests; per-decision robustness flag from the X-Ray; fails closed when no rule certifies an outcome

3 · Wolfsberg AI/ML Principles (2022) mapping

Element	Mechanism
Legitimate purpose	scope bound to FCC alert adjudication under the institution's own versioned policy

Proportionate use	AI agents only investigate and classify; the deterministic engine bounds every outcome; no autonomous closure
Design & technical expertise	open, readable rules; published data contracts; test suite as executable specification
Accountability & oversight	named maker and checker sealed into every record; policy log attributes every rule change to its approver with the replay evidence
Openness	every decision cites its evidence and policy clauses; counterfactual X-Ray available per decision; this pack is reproducible

4 · Policy governance extract

- **Version in force:** v19
- parameter `name_clear_min_mismatches` = 1
- parameter `structuring_low` = 9000
- parameter `structuring_high` = 10000
- parameter `structuring_min_count` = 3
- parameter `mule_min_inbound` = 8
- parameter `mule_max_single` = 600
- parameter `mule_outflow_ratio` = 0.8
- parameter `legal_person_young_since` = 2024
- parameter `tf_overprice_multiple` = 3.0
- parameter `name_similarity_floor` = 0.85
- **Clauses in force:** AML-SC-03, AML-TM-07, AML-TM-09, AML-TM-11, AML-CDD-04, AML-TM-12
- **Behavioural baseline snapshot:** bl-2026-07 — 8 peer groups over 214 cases, mode: advisory (sha `f75c24504d1e3a40...`)
- **Policy log:** 0 activation event(s) — empty policy log

5 · Ledger integrity attestation

- Chain verification: **INTACT** — empty ledger
- Method: every record's SHA-256 seal covers its canonical content plus the previous record's seal; verification re-derives every seal from the genesis record.

6 · Test evidence (enumerated from the suite)

58 tests in CI; claims stated by the suite itself:

- `test_baselines.py::test_advisory_boundary_dispositions_unchanged` — The layer is context: same clauses, same actions, same QA as before it.

- `test_pipeline.py::test_meridian_structuring_goes_to_sar` — TM-only structuring: no screening hit ⇒ suspicious activity → SAR evaluation.
- `test_pipeline.py::test_whitfield_fails_closed` — Nothing fires ⇒ the engine refuses to certify; QA blocks; human review.
- `test_pipeline.py::test_engine_is_conservative_without_rules` — No fired rule ⇒ the engine refuses to certify (routes to a human).

7 · Change inventory (source control)

- 5093c99 · 2026-07-07 · Demo runbook v3: modular 7-to-20-minute flow + demo-beacon.cmd launcher
 - e475cee · 2026-07-07 · Name similarity + trade-mispricing typology + MI report (P14, TM-12, MI)
 - cb362cd · 2026-07-07 · Case clocks: regulatory deadlines as a view over the sealed ledger (P13)
 - ac34eaa · 2026-07-07 · Perpetual Watch: list-delta re-screening + sealed risk re-tiering (P12)
 - 4a01648 · 2026-07-07 · Explainable behavioural baselines - advisory anomaly layer (P11)
 - a655ca5 · 2026-07-07 · Beacon C2: sealed signal exchange between two instances (built + verified)
 - 2575fc5 · 2026-07-07 · Typology packs TM-11 + CDD-04 and jurisdiction profiles v1 (P6 + P10)
 - 24d90c2 · 2026-07-07 · Transparent CRR + seats with four-eyes and compartments (P5 + P8a + P9a)
 - 854f35e · 2026-07-07 · Runbook: full-loop demo beats (Tower, X-Ray, STR export, compliance pack)
 - db9ead1 · 2026-07-06 · goAML STR export with cited narrative + CAdES digest seam (P2, DE/PL)
 - ca7825c · 2026-07-06 · Compliance Pack Generator (P7, EU variant)
 - 1719016 · 2026-07-06 · Decision X-Ray (P3) + Control Tower + integration surface
-

Manifest (generation time, body hash) accompanies this document; regenerate and compare hashes to verify nothing was altered.

```
manifest · body sha256: 703aa88ba203506bc8c7530a30e9cefe9e3d0fd0a662d52025a7c73ffbd62a54
policy v19 · engine 0.1.0 · commit 5093c99 · tests enumerated: 58 · ledger intact: true
```