RISK OPERATIONS · EUROPEAN PAYMENT BUSINESSES · 2026

# The Risk Ops Framework

Detection, investigation, dispute resolution, and compliance documentation as an integrated closed-loop system — not five separate tools.

| 5 stages | 50 days | 1 data model | <45 min |
|---|---|---|---|
| Closed-loop framework | Shadow proof period | Unified across all stages | Regulatory export SLA |

## The core problem with how risk ops is run today

Most payment businesses run risk operations as five separate functions: a fraud model, a rules engine, an analyst queue, a chargeback team, and a compliance function. Each is optimised independently. None shares data. Evidence assembled at Stage 1 (detection) is not available at Stage 5 (dispute resolution). The result is a system that is locally efficient and globally broken.

- **Signal loss:** Fraud detected at the transaction layer is never correlated with chargeback outcomes. The model never learns whether its decisions were correct.
- **Analyst overhead:** Manual review queues grow as false positive rates increase. Analysts spend time triaging, not investigating.
- **Representment failure:** When disputes arrive, the evidence that would have won the case was never retained — or exists in a format that cannot be exported quickly enough for card scheme deadlines.
- **Compliance exposure:** EU AI Act Article 13 requires per-decision explainability logs. Fragmented point solutions cannot produce them.

The closed-loop alternative is not a new tool. It is a different architecture — one where detection, investigation, dispute resolution, and compliance share a single data model from the moment a transaction is processed.

## The 5-Stage Closed-Loop Risk Ops Framework

1

### Stage 1: Real-Time Detection

<100ms resp

- Transaction scored against 500+ signals — behavioural, device, network, identity graph, and historical pattern
- Explainability log written at decision time: signal weights, score, model version, decision rationale
- Human-readable rationale generated for every decision — not just declines
- Linked-account graph queried for cross-customer fraud ring patterns that point solutions miss
- Low-confidence decisions auto-routed to analyst queue (typically ~20% of flagged volume)

## 2 Stage 2: Analyst Investigation

Structured qu

- Flagged transactions arrive in analyst queue pre-assembled with full evidence packet
- Evidence includes: transaction data, identity signals, device fingerprint, linked accounts, and Stage 1 rationale
- Analyst reviews the case — not the raw data. Decision: confirm fraud / clear / escalate
- Analyst outcome feeds back into the model as a labelled training signal
- Queue prioritised by risk score, time sensitivity, and dispute deadline proximity

## 3 Stage 3: Automated Dispute Resolution

Chargeback e

- Chargeback evidence packet assembled automatically at Stage 1 — not retrospectively when dispute arrives
- Evidence includes all Stage 1 signals, Stage 2 analyst outcome, and EU AI Act explainability log
- Representment package formatted for Visa/Mastercard dispute workflows
- VAMP threshold (0.9%) monitored continuously — alert triggered before acquirer review threshold reached
- Win rate benchmark: 60–75% representment success rate across ICP transaction profiles (OfferSpine v4 benchmark — not a guarantee)

## 4 Stage 4: Compliance Documentation

EU AI Act Art

- Full per-decision explainability log exportable on demand — target: under 45 minutes
- Log is human-readable, timestamped, and immutable
- Covers all automated decisions — approvals, declines, and escalations
- Formatted for regulatory submission to FCA, BaFin, DNB, MGA, UKGC, and equivalent authorities
- August 2026 EU AI Act enforcement: shadow mode builds the baseline audit trail during the proof period

## 5 Stage 5: Closed-Loop Learning

Analyst feedb

- Stage 2 analyst outcomes (correct/incorrect decisions) feed back into the model as labelled training data
- Model trained on your specific fraud profile — not generic industry benchmarks
- Drift detection monitors for changes in fraud pattern distribution that trigger retraining
- VAMP trajectory and chargeback rate delta tracked as lagging confirmation of detection improvement
- Day 50 shadow mode report: first baseline measurement of closed-loop improvement potential on your data

## The unified data model: why it matters

The framework only works if all five stages share the same data model. Evidence assembled at Stage 1 must be available — without transformation or reformatting — at Stage 5. This single constraint rules out most point-solution stacks.

| Point-solution stack | Closed-loop architecture |
|---|---|
| • Stage 1 detection: fraud model output | • Stage 1 detection: full evidence packet written |
| • Stage 2: CSV export to analyst tool | • Stage 2: evidence pre-assembled, analyst reviews |
| • Stage 3: manual evidence assembly from 3 systems | • Stage 3: representment packet auto-generated |
| • Stage 4: impossible without full log | • Stage 4: exportable in <45 min, regulator-ready |
| • Stage 5: no feedback loop | • Stage 5: analyst outcomes retrain the model |

## Implementing the framework: the shadow mode path

The fastest way to validate whether a closed-loop architecture would improve your detection delta is to run it in parallel — without replacing your existing stack. SENTR's 50-day shadow mode evaluation is designed for exactly this.

- **Week 1–1.5:** Architecture Session + 8-day API integration. Read-only connection to your transaction and chargeback data.
- **Days 1–50:** SENTR runs in shadow mode. Your stack makes every live decision. SENTR observes, scores, logs in parallel.
- **Day 50:** Intelligence report delivered. Detection delta quantified on your actual transaction volume. EU AI Act compliance gap assessed. Board-ready ROI snapshot produced.
- **Post-Day 50:** You decide. No invoice until you've seen the data. No commitment until the evidence is in front of you.

---

## Book your Architecture Session

The 20-minute Architecture Session is where we map your current risk ops stack against the closed-loop framework and show you exactly where the detection gaps are. No pitch deck. No commitment.

**Book at sentr.io/proof**

hello@sentr.io · sentr.io · EU AI Act compliant · GDPR compliant · All data processed within EU

Win rate benchmark (60–75%) is an OfferSpine v4 benchmark range — not a guarantee. Your actual results depend on transaction mix, fraud profile, and dispute filing practices.