Ledgerproof

Turn one data export into a quantified risk diagnosis.

Diagnostic-led risk & compliance for regulated gambling. The free Diagnostic proves the problem; the paid Blueprint explains the fix.

Multi-jurisdiction. Benchmarked against UKGC RTS as a high-control reference framework.

NOT AN OPERATOR · HOLDS NO FUNDS · ACCEPTS NO BETS

Advanced IT Consulting (UK) · Confidential

Upload → Diagnose → Quantify → Trust → Implement.

01

Upload

Send one anonymised export. No integration, no new instrumentation.

02

Diagnose

The engine determines what your data can prove — and surfaces findings only where the evidence supports them.

03

Quantify

Where evidence supports it, measurable commercial leakage is sized: affected area, estimated recoverable range, and confidence.

04

Trust

Every figure traces to observed data; every decision is explainable and audit-ready. Nothing is fabricated.

05

Implement

The paid Implementation Blueprint explains exactly how to fix it; Integrated Deployment captures the value in production.

The free Diagnostic proves and quantifies the problem where the evidence supports it — and never fabricates a finding the data cannot support.

Free proves the problem. Paid explains the fix. Deployment captures the value.

Free Diagnostic FREE

Proves the problem and quantifies what is at stake.

· The category of each problem found (where evidence supports it)

· Its financial magnitude and estimated recoverable range

· The affected area and the owning business domain

· A confidence level on every figure

Implementation Blueprint PAID

Explains exactly how to eliminate it.

· The remediation design for each opportunity

· The specific control and rule changes

· The rollout sequence and configuration values

· Success criteria and verification

Integrated Deployment ENGAGEMENT

Captures the value in production.

· Live integration on your stack

· Staged adoption: shadow → advisory → controlled automation

· Ongoing audit evidence and monitoring

· Owned by your team, with Ledgerproof support

The Diagnostic shows a problem's category, magnitude, affected area, confidence, ownership and estimated recoverable range — never the thresholds, rules or steps that resolve it. That boundary is the product contract.

The leak is real — quantified on your data, not assumed.

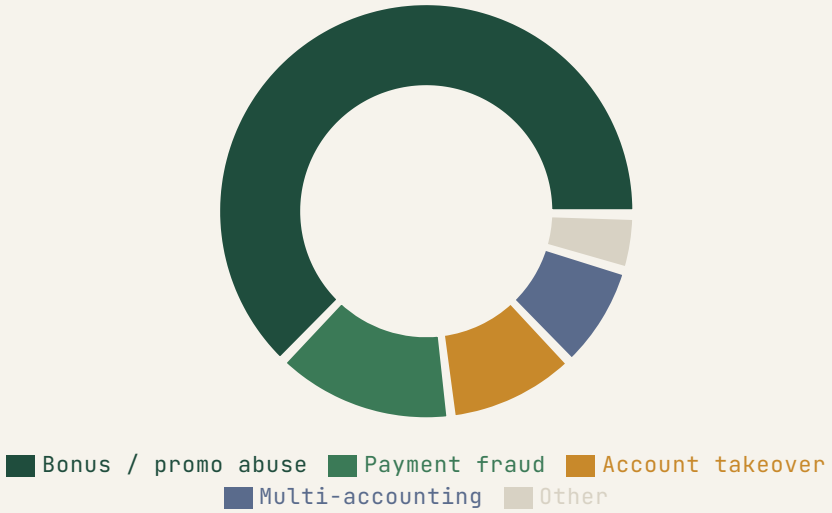

Fraud composition — Sumsub, 2025 iGaming Fraud Report. Industry context; your figure is quantified from your own data where the evidence supports it.

~64%
of iGaming fraud is bonus / promo abuse
According to Sumsub 2025.

10–20%
of promotional turnover lost to abuse
According to LexisNexis.

Every false decline
= lost player lifetime value
According to industry.

Regulators just made this expensive.

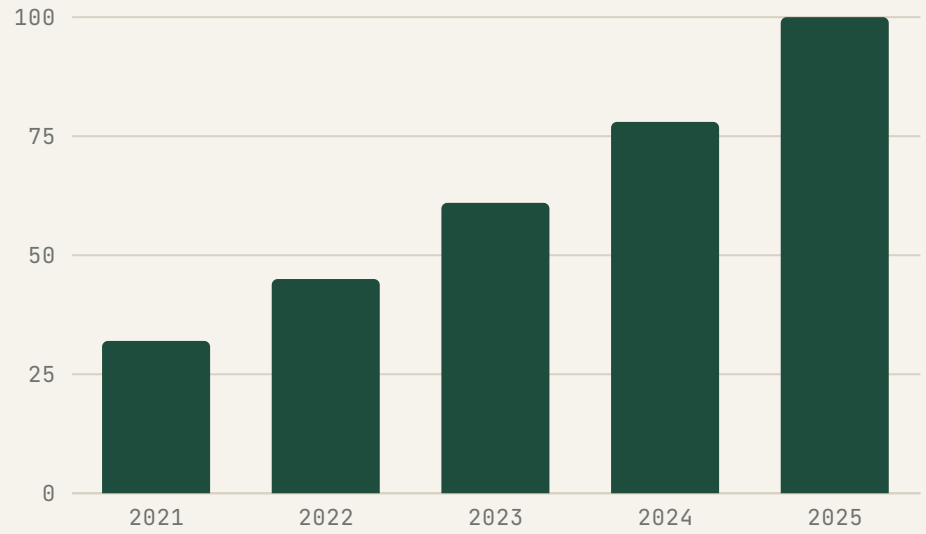

Indexed enforcement scale — illustrative, from public UKGC enforcement summaries.

- Regulators named alert backlogs a mechanism of failure.
- Static thresholds are penalised — the Betfred source-of-funds case, an example from one strict regime.
- RTS tightened: the audit trail now ranks alongside AML.

Risk-based catches more abuse and blocks fewer good players.

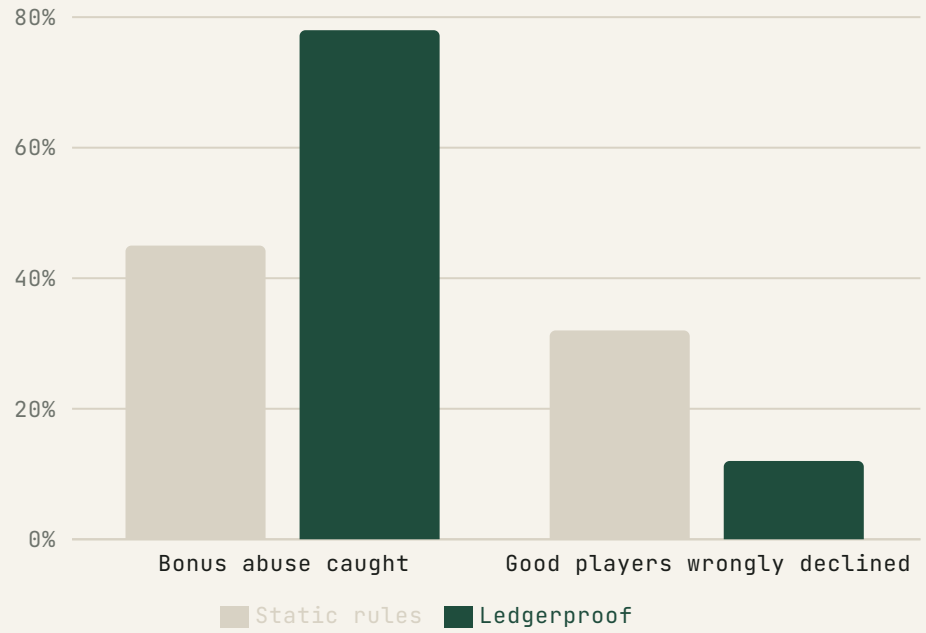

Illustrative, based on synthetic simulation. Your figures are measured on your data during the pilot.

The gap between static and risk-based decisions — same data — is the measurable leakage the Diagnostic quantifies, where the evidence supports it.

Two policies, same data — the gap is the opportunity.

Proof in five clicks — and proven on real data.

1. **Decision**
   accept / reduce / reject / review / cooling-off
2. **Rule + version**
   which policy fired, at which version
3. **Input data**
   the exact inputs behind it
4. **Owner + time**
   accountable owner, timestamp
5. **Tamper-evidence**
   hash chaining the record

audit-readiness, not a passed audit. Evidence for an independent auditor.

Any decision resolves to its rule, version, inputs, owner, and a tamper-evident hash — a trail an auditor can follow.

We ran the engine on a real public dataset (Mendeley 9j5gcygnwg, CC BY 4.0). On payments-only data it produced zero findings — it refuses to fabricate what the data cannot evidence. See the Validation page.

What it's worth — and how the value is captured.

Send one export.

Ledgerproof determines what your data can prove, identifies measurable commercial leakage where evidence supports it, and quantifies the value at stake.

Full remediation logic is delivered through the paid Implementation Blueprint. We do not promise a number in advance.

Blueprint → Integration

The paid Implementation Blueprint explains exactly how to fix each opportunity — remediation design, control and rule changes, rollout and verification, proven in a Shadow-Mode pilot. Integrated Deployment then captures the value in production.

The Diagnostic is free and read-only. Nothing in production changes until you decide to act.

The credibility: an explainable decision engine.

01
Decide
signed, timestamped record

02
Reserve
signed, timestamped record

03
Confirm
signed, timestamped record

04
Audit
signed, timestamped record

ACCEPT REDUCE REVIEW REJECT COOLING-OFF

Every decision is defensible.

bet	£5,000
bankroll	£1,000,000
player risk	0.15
open liability	£40,000

decideRisk()
deterministic

ACCEPT

01 · within risk-adjusted capacity
02 · utilization 0.23 < 0.85
03 · player-risk below review threshold

Benchmarked against UKGC RTS — and we show our gaps.

RTS 1	Customer account information	PARTIAL
RTS 5	Result determination	COVERED
RTS 10	Interrupted gambling	COVERED
RTS 12	Financial limits	BOUNDARY
RTS 14	Responsible product design	COVERED
Security §4	ISO/IEC 27001:2022 Annex A	COVERED

Legend: COVERED PARTIAL BOUNDARY GAP N/A

RTS is our chosen benchmark — a high-control reference framework, with jurisdiction-specific mappings for Curaçao, MGA, Anjouan, Estonia and other regulated markets. We show the gaps (RTS 1 retention, RTS 12 = operator/PAM boundary, 14B, 8.17) on purpose — an honest map strengthens an auditor's trust.

Everyone owns one layer. We connect all three.

KYC / AML

signals, no decision

RG analytics (Neccton, Mindway)

harm view, no money / crash exposure

Payment gateways

money, but blunt risk

Ledgerproof

decision + justification + audit — the connective layer

Built, tested, and honest about its limits.

820	6.3M	~353k	17
TESTS PASSING	ROWS LOAD-TESTED	TX / SEC	RTS MAPPED

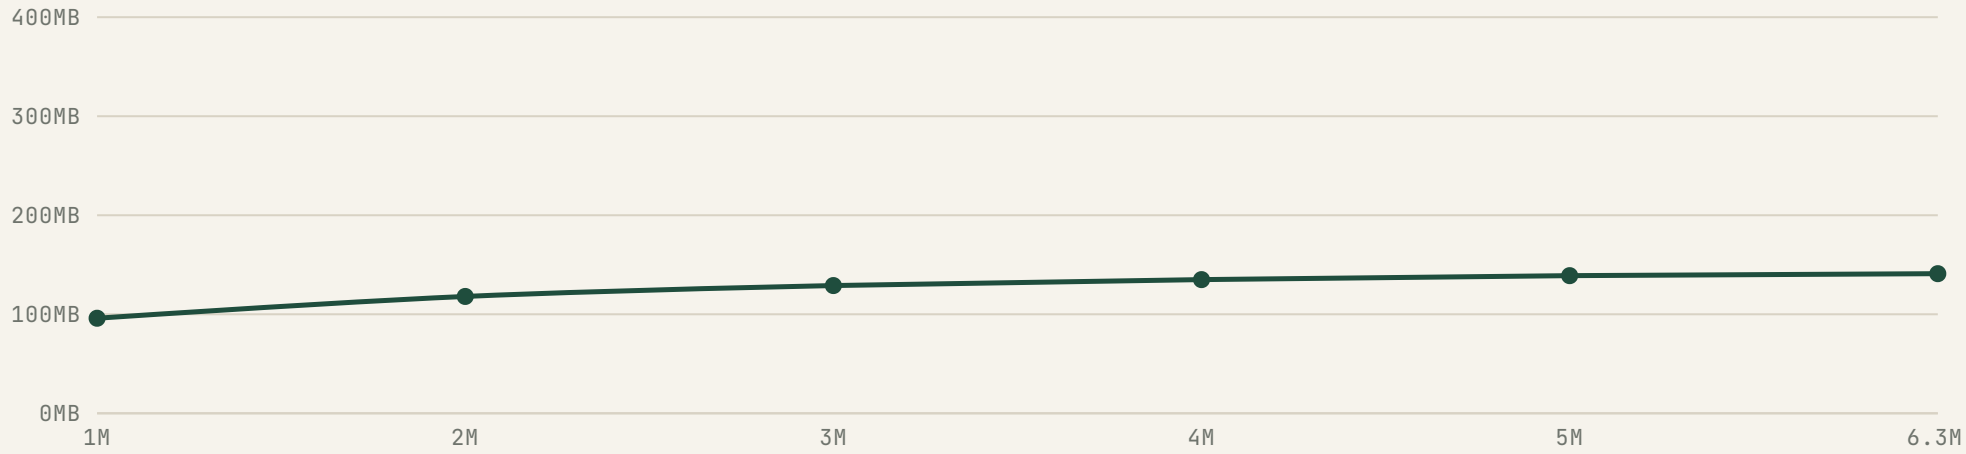

Load test on synthetic PaySim data (mobile-money, not iGaming); performance only, not a revenue claim. Review-candidate: no money, custody, or settlement. Flat memory demonstrates O(1) streaming under load.

Ledgerproof

Send one export. See what your data can prove.

Request a pilot

Advanced IT Consulting (UK) · pilots@ledgerproof.example