



AML · TRANSACTION MONITORING · REGULATORY REPORTING · SECURITIES SURVEILLANCE

Compliance you can prove to the regulator.

One platform from the first transaction to the filed goAML report — with AI that drafts, and humans who decide.

Banks

SACCOs

Digital lenders

PSPs & fintechs

Investment banks & brokers

THE PROBLEM

Compliance teams are drowning — and still accountable.

Regulators expect timely, accurate, explainable reports. The tooling most institutions run makes that harder, not easier.



Alert fatigue

Rules tuned once and never revisited produce mountains of false positives. Analysts triage noise instead of investigating risk.



Disconnected tools

Monitoring, screening, case management and filing live in separate systems. Data is re-keyed, evidence is lost, and nobody can reproduce a decision.



Filing under pressure

An STR is assembled by hand from several systems, validated only when the regulator rejects it, and a schema change means a rebuild.



AI without guardrails

Leaders want AI's speed, but cannot let it touch regulated decisions, leak customer data, or act without an audit trail.



Examiner scrutiny

"Who approved this, on what evidence, and can you reproduce it?" is answered from spreadsheets and memory.



New asset classes

Investment banks and brokers need market-abuse and securities-transfer surveillance on top of AML — usually a separate product.

THE ANSWER

One accountable platform, end to end.

Detect, investigate, decide, report and evidence — in a single system built around the people who answer to the regulator.



■ System-driven ■ Human decision ■ AI-assisted draft



Everything a compliance function needs, in one place.



Ingestion

- Signed, idempotent REST API
- Batch files with quarantine
- ISO 20022 · M-Pesa · FIX
- No-code Mapping Studio



Monitoring

- Real-time rule engine
- Visual + plain-English builder
- Backtest before go-live
- Network / mule detection



Screening

- Sanctions · PEP · adverse media
- Fuzzy name matching
- Four-eyes dispositions
- Periodic full-book re-screen



Risk & EDD

- Reproducible customer risk rating
- Auto-opened EDD cases
- EWRA · peer segmentation
- Securities risk factors



Cases

- Workflow with SLA timers
- Evidence & timeline
- Link-analysis graph
- STR tipping-off controls



Sentry Agent

- Cited investigation memos
- Read-only, PII-redacted
- Three model tiers
- Per-feature kill switch



goAML reporting

- STR · SAR · CTR
- 3-layer validation incl. XSD
- FIU rejection → correction loop
- Deadlines & registers



Analytics

- Executive dashboards & pivots
- Metric Designer with alerts
- Typology Radar
- Securities & reporting KPIs

DETECTION

Rules your team can see, tune and prove.

Every rule is versioned data with named, tenant-overridable parameters — not code. Change it, backtest it on your own history, then activate it with maker-checker.

32 scenarios out of the box

Mapped to recognised typologies, each with a backtest fixture and coverage-map entry.

Structuring · 7

Layering · 8

Money mule · 2

Cash-intensive · 3

Turnover · 2

Velocity · 3

Fraud · 2

Market abuse · 5

Author it your way

- Visual rule builder, round-trip with the JSON rule definition
- Describe a rule in plain English — AI drafts it, backtest runs automatically
- Rule performance dashboard flags noisy rules
- Rich aggregates: counts, sums, distinct counts, ratios, netted volume

Change a rule safely

Draft

Edit parameters or logic

ANALYST

>

Backtest

Projected volume & false-positive estimate

SYSTEM

>

Approve

Author ≠ approver, step-up MFA

MLRO

>


Activate

Engine loads only Active versions

SYSTEM

Faster investigations. Humans keep the decision.

The Sentry Agent drafts; people act. AI has no write path to anything with regulatory consequence — by design, not by policy.



PROCESS FLOW

From alert to filed report — every step owned and evidenced.



■ System ■ Human, separation-of-duties enforced

Reports that are right before they leave the building.



Built from your records

Subject, accounts, transactions and indicators are assembled from customer, account and transaction data — with a provenance map showing what came from where and what is missing.



Three-layer validation

- Data completeness
- Business rules & code tables
- XSD — errors pinpointed to field and XML line, with a fix hint



The rejection loop, closed

Paste the FIU response; reasons are parsed; a correction is drafted as v2 with the original reference; a different MLRO approves; the original is superseded.



Schema as content

Versioned schema packs: a new regulator schema is installed, mapped and regression-checked — not re-engineered. Filed reports keep the schema they were built under.



Deadlines that track themselves

On-track, due-soon and overdue per configured jurisdiction timetable; the clock restarts at an FIU rejection. Automatic daily CTR generation.



Examiner-ready registers

STR, CTR, declined-business and training registers; board packs; every view and download of an XML logged with its checksum.

Output today is built on a generic goAML-style schema and labelled as such; the regulator's published schema installs as a pack. Submission to the FIU is a controlled hand-off unless an API/SFTP channel is available.

AML and securities surveillance, one platform.

Orders, trades and free-of-payment transfers are first-class events — so market-abuse and custody patterns sit beside AML typologies, not in another product.

Intake

- **FIX drop-copy** (new order, cancel, execution)
- REST and blotter files; instrument master
- Idempotent; trades reference customers by your own account ids

Customer risk

- Instrument risk (derivatives, unlisted) and venue risk (OTC share) feed the rating

Seven surveillance scenarios

Wash / matched trading

Layering / spoofing

Order-to-trade ratio

Concentrated accumulation

Venue fragmentation

Free-of-payment transfers

Custody pass-through

Surveillance dashboard

Orders, executed trades, cancel rate, alerts by scenario, value by instrument and venue — from a dedicated analytics store.

Insider-dealing, front-running and marking-the-close detection require announcement and market-close data; they are on the roadmap, not in today's scope.

USER GROUPS

The right tools for every role. Nothing more.

GROUP	WORKS IN	CAN	CANNOT
Analyst / Senior Analyst	Alerts · Cases · Customer 360 · Screening · Rules	Investigate, build cases, request AI memos, propose dispositions and STRs, backtest rules	Approve their own work; see restricted STR data
MLRO / Deputy MLRO	Approvals · Reporting · KPIs	Confirm and file STR/SAR/CTR, record rejections, approve corrections, release batches	Approve their own corrections (a second MLRO must)
Risk Manager	Risk · EDD	Tune and score customer risk, run EDD	File reports
Administrator	Admin · Tenants · Ops console	Configure tenants, users, reference data, schema packs, jobs and health	File reports or approve regulated decisions
Auditor / Regulator	Audit trail · Sample reviews	Read everything they are entitled to; record review findings	Change anything — every action is disabled
Integration engineer	API · Mapping Studio	Connect and monitor source systems	Approve or file regulated reports

Built so the answer to "prove it" is one click.



Maker-checker, enforced

Separation of duties is checked in the handlers, not only the screens: the preparer can never be the approver for STR corrections, rule activation, screening confirmations or CTR release.



Strong authentication

TOTP and passkeys with step-up for sensitive actions; single sign-on ready. Default-deny authorisation on every page and endpoint.



Tamper-evident audit

Who, what, before and after — hash-chained so alteration is detectable, with database-level immutability for the application role.



Tenant isolation

Tenant id on every tenant-owned table with enforced query filters, covered by isolation tests; optional row-level security as defence in depth.



Tipping-off control

STR-linked cases, report XML and even reporting aggregates are visible only to the privileged policy; every export is watermarked and logged.



Data protection

No PII in logs; PII pseudonymised before any AI call; retention policies with signed deletion certificates; signed webhooks and HMAC-verified ingestion.

See the whole programme — without touching production data.

Dashboards read purpose-built analytics stores fed by events, so heavy analysis never slows monitoring.

Operations

- Executive dashboard & role workspaces
- Pivot analytics with Excel export
- Case productivity, SLA compliance, QA sampling

Intelligence

- Typology Radar: emerging-pattern detection
- Period-over-period league tables
- Metric Designer with RAG status and breach webhooks

Regulatory & securities KPIs

- Open / overdue filings, rejection rate and reasons
- Orders, trades, alerts by scenario, value by instrument
- Signed embeds for client portals

Events

Outbox-published facts

Facts

Insert-only, idempotent

Rollups

Columnstore, rebuildable

Widgets

Role-gated, tenant-scoped

TECHNICAL OVERVIEW

A modular platform, engineered for regulated workloads.

EXPERIENCE & ACCESS

Blazor Server UI

Minimal APIs (ingestion · egress · webhooks)

OIDC SSO · MFA · passkeys

DOMAIN MODULES — ONE SCHEMA, ONE DBCONTEXT EACH

Ingestion

Monitoring

Screening

Risk

Cases

Intelligence (AI/ML)

Reporting

Analytics

Identity

Admin

CROSS-CUTTING

Transactional outbox + event bus

Background job framework (leader-elected)

Hash-chained audit

Multi-tenancy & query filters

OpenTelemetry · structured logs

DATA

SQL Server (temporal tables, columnstore)

Redis cache / backplane

Vector store for retrieval

ONNX models in-process

Enforced boundaries

Modules talk only through contracts and events; architecture tests fail the build if a boundary is crossed.

Deploy anywhere

.NET 10; IIS or container; rolling, additive-first migrations; documented DR runbook.


Open integration

HMAC-signed APIs, ISO 20022, FIX, webhooks, field-scoped egress, signed embeds.

GETTING STARTED

From signature to first live filing.

A structured, owner-by-owner plan with entry and exit criteria at every step. Indicative 10–16 weeks, dominated by data quality and tuning — not software.



Timelines and service levels are indicative and subject to commercial agreement.

What sets it **apart.**

1 · Filing is part of the product, not an afterthought

Detection, investigation and the validated regulatory report live in one system with one audit trail — no re-keying between tools.

2 · AI with a hard boundary

Useful drafts, cited evidence, full logging — and no technical path to a regulated decision. Compliance leaders can adopt it without arguing with their examiner.

3 · Provable accountability

Separation of duties enforced in code, hash-chained audit, reproducible risk scores and reproducible alerts. "Prove it" has an answer.

4 · Future-proof reporting

Regulator schema changes arrive as content packs with a compatibility harness — not as a rebuild — and filed history is never rewritten.

5 · AML and securities surveillance together

Banks, SACCOs, fintechs and investment houses on one engine, one data model and one governance model.

6 · Configurable, not customised

Thresholds, lists, queues, workflows and deadlines are tenant settings. Every feature is built to be sold to the next client without a code change.

What is live, and what is next.

Available today

- Ingestion, monitoring, screening, risk, cases
- STR · SAR · CTR with XSD validation and rejection loop
- Sentry Agent and ML alert scorer
- Securities scenarios, FIX intake, dashboards
- Maker-checker, step-up MFA, hash-chained audit

On engagement

- Install the regulator's official schema pack and code tables
- Additional report types (e.g. EFT / IFT) on the official schema
- Client-specific source connectors
- Tuning on the client's own history

Roadmap

- Insider-dealing and front-running detection (needs announcement data)
- Single reporting-currency normalisation across all thresholds
- Self-service training-record entry
- Additional SSO integrations

Roadmap items are directional and not a commitment.



NEXT STEPS

See it on **your data.**

A scoped proof of value: load a sample of your customers and transactions, tune three scenarios on your history, and walk an STR from alert to validated report with your MLRO in the room.

1 · Discovery call

Sources, report types, jurisdiction.

2 · Proof of value

Your data, your scenarios, your MLRO.

3 · Plan & proposal

Timeline, owners, service model.