AMANAH

SOVEREIGN AI INFRASTRUCTURE

From Alert Overload to Actionable Intelligence.

Amanah AML

Rules, machine learning, network detection and an investigation agent, as one engine, on your state, under your law.

Six of the ten largest already run this.

Turning enterprise data into trusted regulatory intelligence.
Not a research demo, and not a product we are learning on with you.

01

BUILT FOR BUSINESS OUTCOMES

We are measured on recall, on analyst workload and on cases closed, not on benchmarks. Every capability here exists because a compliance team asked for it.

02

FOUNDED TO CLOSE A GAP

Advanced AI research rarely survives contact with a regulated bank. Amanah carries it across, into systems that pass audit and stay in production.

03

TAILORED, NOT OFF THE SHELF

AI specialists, data scientists and financial crime practitioners, building for your typology mix and your regulator. Not a fork you maintain alone.

TEAM EXPERIENCE SPANS

IBM SAS Google specialist AI firms

Proven where the volumes are largest.
Delivered where your law applies.

700 million transactions later.

In a production deployment at a top five retail bank:
200 million customers, $2.1 trillion in customer assets,
700 million transactions analyzed.

99.6%

recall within the top 70% of ranked cases

The real cases are at the top of the queue, not buried in it.

30%

fewer false positives, and 30% less analyst workload

The same team reviews less and finds more, because the queue is ordered by risk.

200%+

detection gain over rule based alone

Cases the rules structurally could not see, because the evidence was never in one row.

98%

match to senior AML analyst judgment

Every alert carries the read your best reviewer would give it, at machine speed.

Four capabilities. One engine.

Called together on every alert, through one orchestration layer, on one audit trail. Each one stands alone, and each one is measured.

01

AI powered alert prioritization

99.6% RECALL . 98% ANALYST MATCH

Every alert scored and ranked, so the queue arrives by risk instead of by arrival.

02

Optimized network detection

200%+ DETECTION GAIN . +15% MORE MEMBERS

The ring read as a network, with the expensive computation run only where it changes the answer.

03

Document intelligence

80% HARD CASES SOLVED WITHOUT ANY FINE-TUNING

The risk signal pulled out of stamped, handwritten, out of order paper.

04

Explainable reports

5 / 27 INDICATORS . 100+ METRICS

Every score decomposed, every filing drafted with its evidence, and human gated.

Rules catch the known. The model prioritizes the subtle. The network reveals the hidden.

Alert overload.

A mid size bank raises thousands of alerts a day against fixed capacity: 70 to 85% false positives, and judgment decaying past 250 alerts per analyst. The rules were never enough.

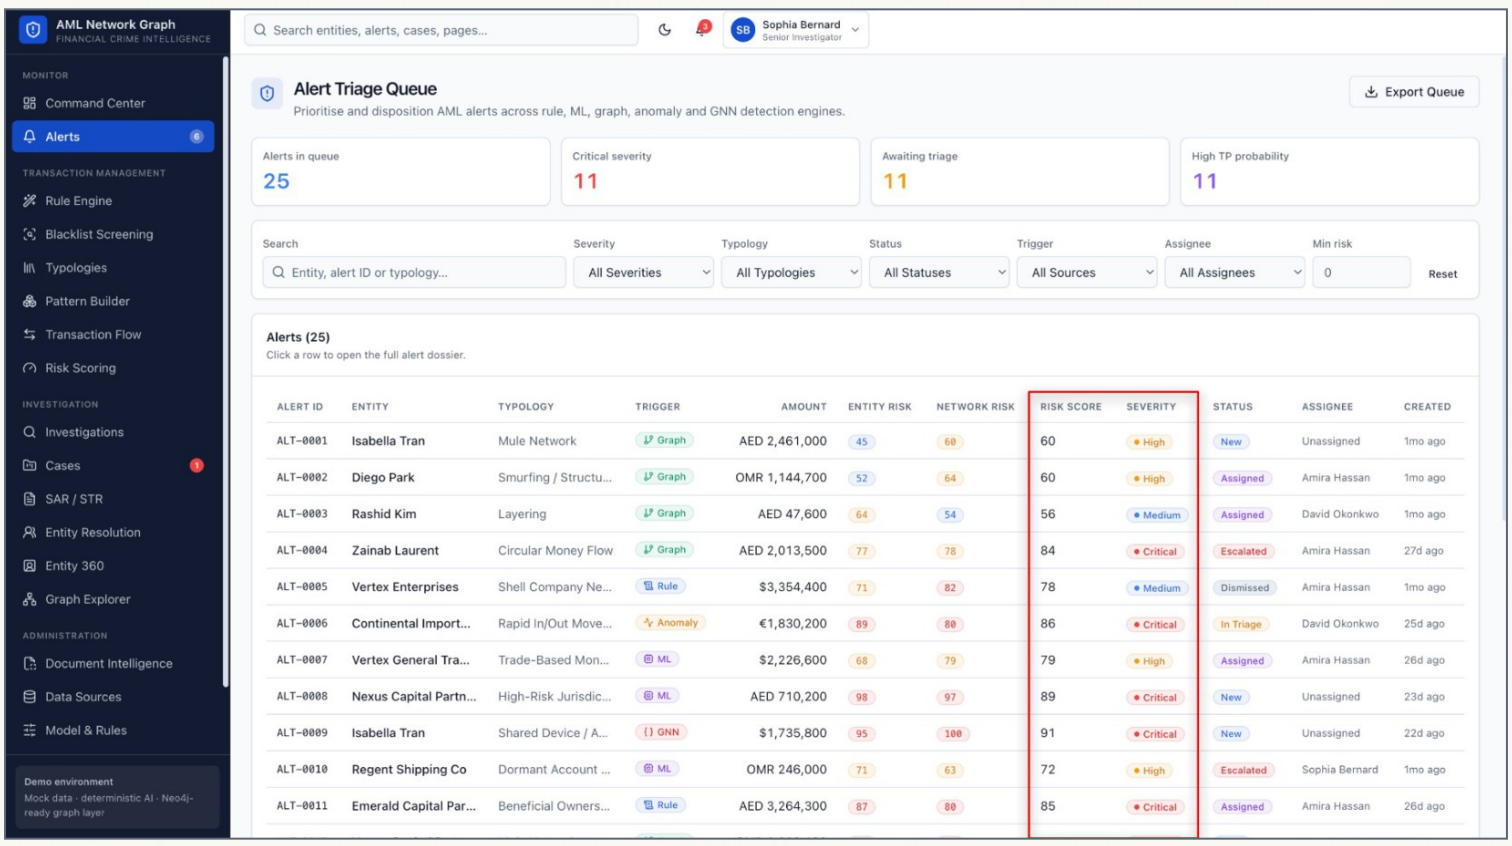

ALERT TRIAGE QUEUE: EVERY ALERT SCORED, THE QUEUE ORDERED BY RISK NOT ARRIVAL.

HOW THE PRODUCT SOLVES IT

The same alerts your rules already raise, reordered so the real cases arrive at the top. Scored on six families of signal, from basic and sequential to customer profile and network. Everything uncertain still reaches a person.

WHAT IT PRODUCED IN AN IMPLEMENTATION

99.6%
recall within the top 70% of ranked cases

98%
match to senior analyst judgment

30%
less analyst review workload

Ordering the queue finds the cases that were already in it. It cannot find the ones that were never in a single row.

Criminals work as networks. Your rules read one row.

Five accounts, every one within limits: $9,400. $9,700. $9,200. $9,850. $9,500. Checked one at a time no rule fires. Seen together they feed one consolidation hub and exit through one shared address.

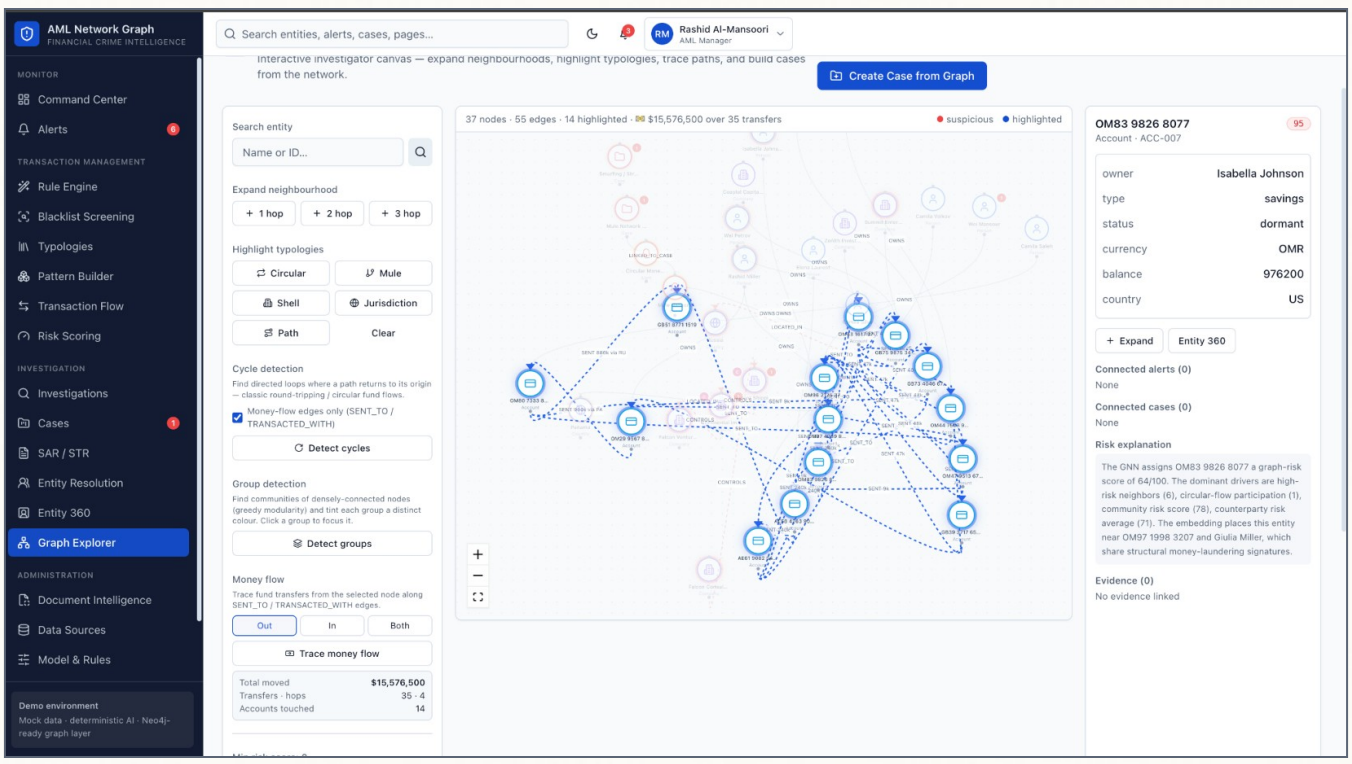

NETWORK VIEW: MONEY FOLLOWED ACROSS HOPS, WITH THE PATH THAT JUSTIFIES THE ALERT.

HOW THE PRODUCT SOLVES IT

An ontology built for financial crime: people, companies, accounts, transactions, addresses, devices and countries. A dense cluster is the ring, convergence on a hub is the consolidation point, a directed chain over time is layering.

WHAT IT PRODUCED IN AN IMPLEMENTATION

200%+

detection gain over rule based alone

+15%

more network members found than rules alone

70 / 80%

network and member level accuracy, measured independently

The network is only as good as what you can put in it, and most of what a bank knows is not in a table.

Most of what you know is not in a table.

It is in their documents, and their documents are not clean scans. Stamped over. Handwritten. Photocopied crooked. Out of order. That is the normal case.

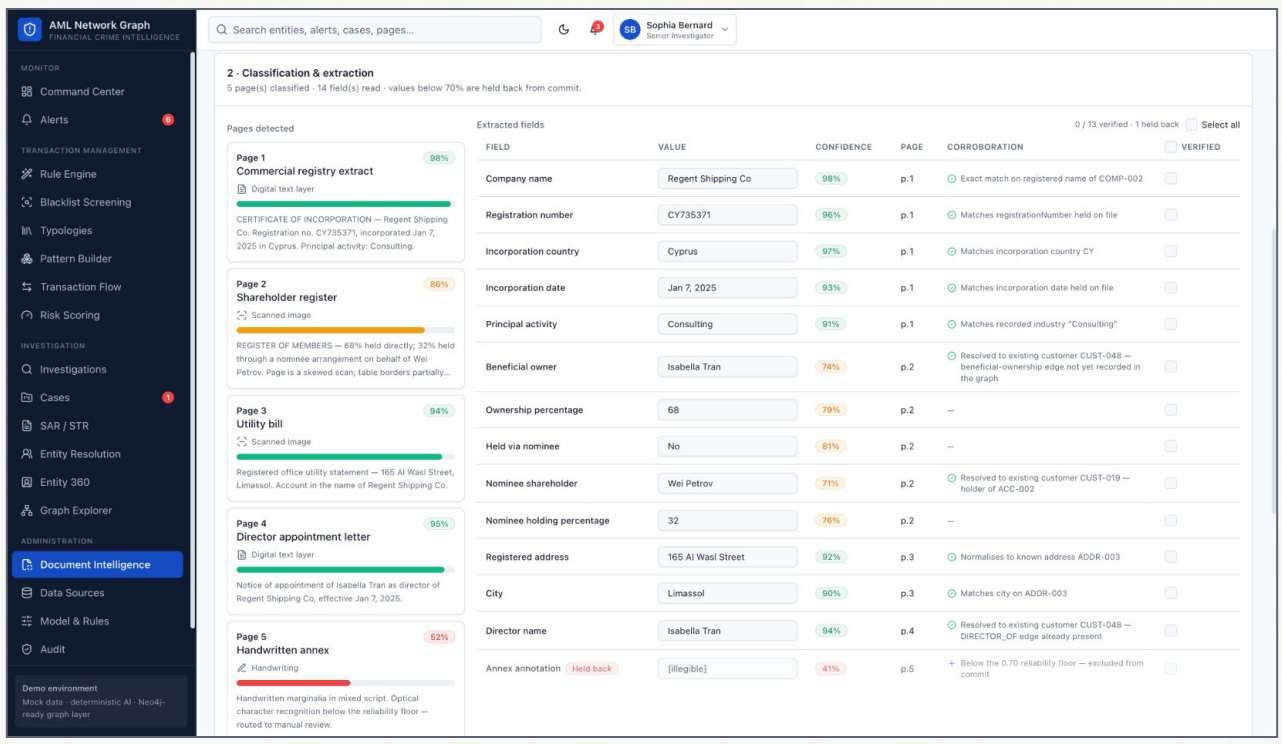

KNOWLEDGE EXTRACTION: EVERY FIELD LINKS BACK TO THE PAGE IT CAME FROM.

HOW THE PRODUCT SOLVES IT

Eight stages, not one: extract, clean, classify, structurize, rules, validate, derive and ingest. Deterministic code does the deciding wherever it can.

WHAT IT PRODUCED IN AN IMPLEMENTATION

80%
of the hard cases solved, no fine tuning, in a proof of concept

40
categories of material identified in one engagement

8
stages, each one auditable on its own

Documents become the entities and relationships the network reasons over.

A model you cannot explain is a model you cannot defend.

Not to us. To your regulator, your auditor,
and eventually to a court. Why did it score that?
Who decided? What exactly am I signing?

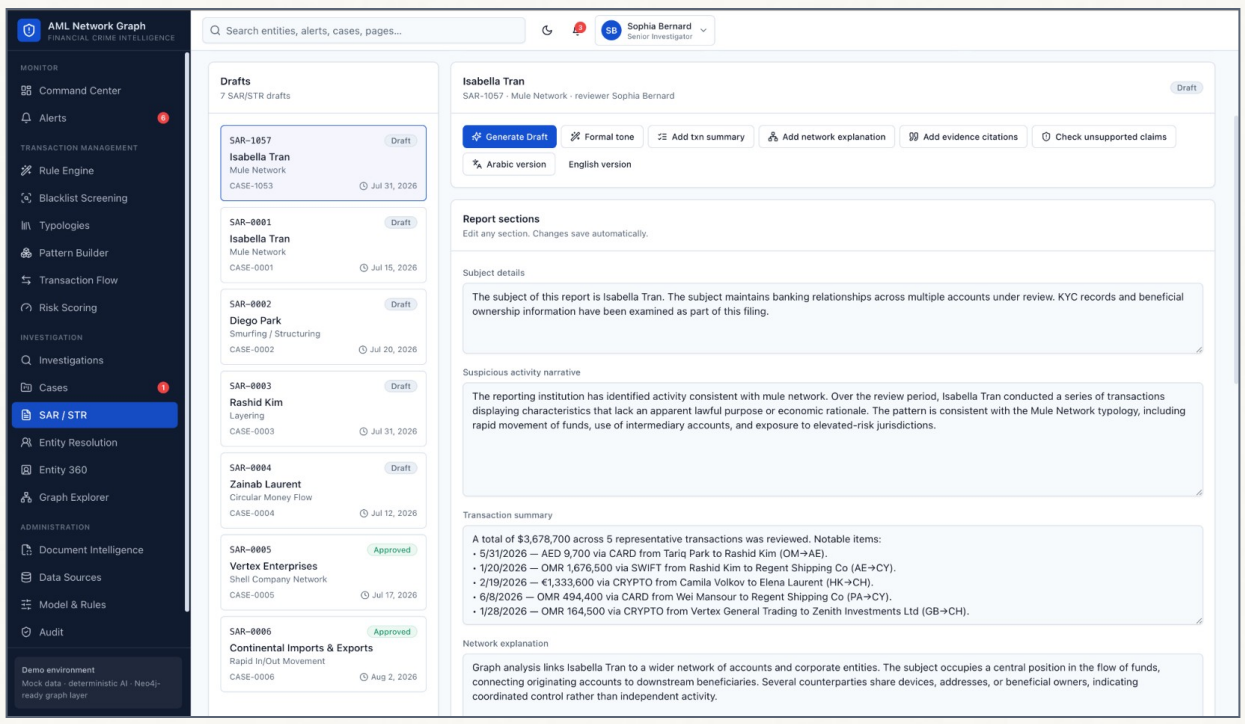

REPORT GENERATION: THE SYSTEM DRAFTS OVER THE NETWORK PATH, CITING ITS EVIDENCE.

HOW THE PRODUCT SOLVES IT

Every score decomposes into the indicators that produced it.
Your reviewer edits, your MLRO approves, your bank files.
No report is ever submitted by this system.

WHAT IT PRODUCED IN AN IMPLEMENTATION

0 to 100
case suspicion score,
indicators named and weighted

5 / 27
indicator categories
and sub categories

2h to 15min
STR drafting time,
a product capability not a
deployment result

Four capabilities, one engine.
The question left is what it costs to run.

Graph analytics is expensive. Ours is sized to the question.

The return does not come from running every algorithm on every node. It comes from running the right algorithm on the right sub graph, and nothing else.

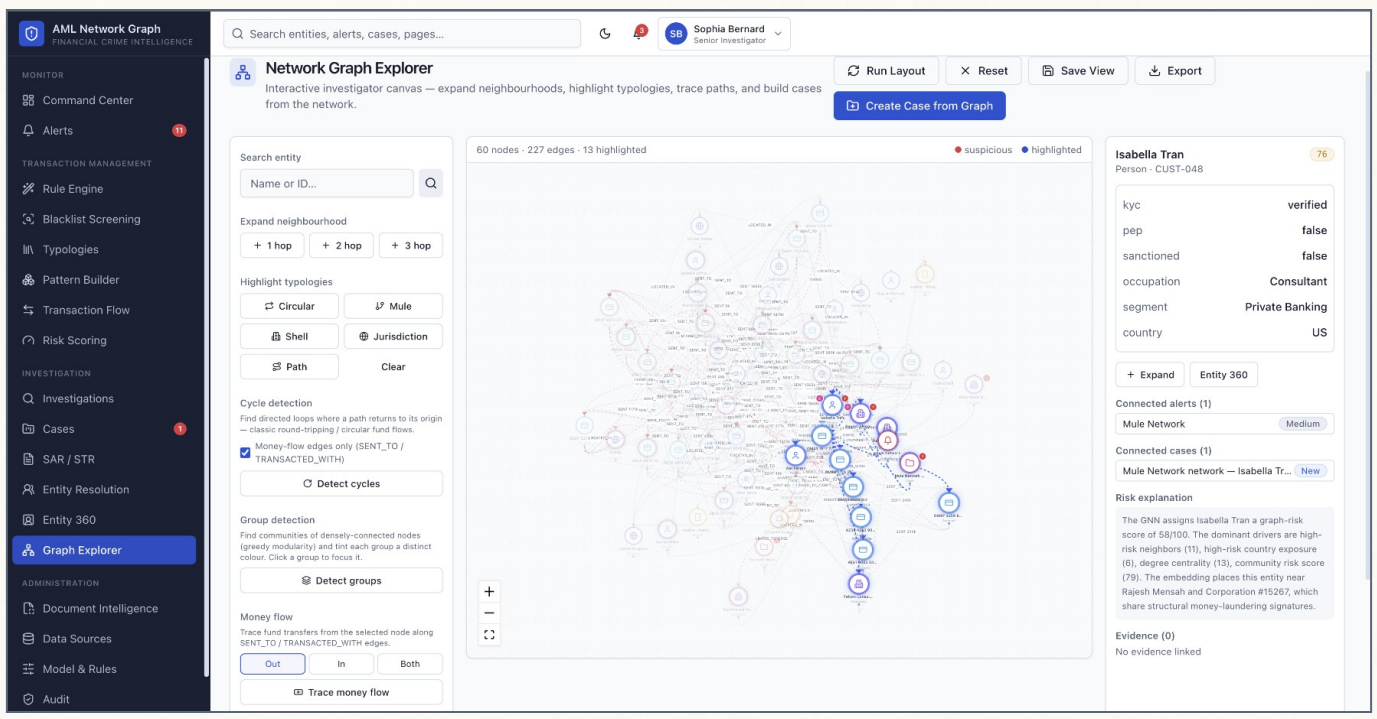

GRAPH EXPLORER: HOP DEPTH AND ALGORITHM CHOSEN PER QUESTION. 60 NODES IN VIEW.

GRAPH IS SIZED	Selected by the question, not loaded whole.
ALGORITHM IS CHOSEN	Community, centrality, cycle or fund flow.
DEPTH IS SET	By the motif hunted, never maxed for itself.
WEIGHTS ARE LEARNED	From confirmed cases, not assigned by hand.

8 hops
Still in seconds.

100M
Nodes, at sub second query.

Under 30s
To load 67 million edges.

Run the expensive computation only where it changes the answer, and network detection is affordable to leave switched on.

The full platform, application to data.

The AML engine sits on one modular stack, end to end. Every layer below is shared, which is why a capability added later joins the same engine rather than arriving as a second system.

APPLICATION	Intelligent AML			
AI DEVELOPMENT	AI Developing Suite	Graph Application Suite	Agent Developing Suite	
OPERATORS AND MODELS	ML / DL	Graph	LLM	Specialized AI
DATA	Data Governance	Graph Database	Real time Computing	Labeling

One application on this page, deliberately. The same stack carries risk control, marketing, service and operations, and a companion document AI stack follows the same four layer shape.

One stack. Which means the second use case costs a fraction of the first.

Where we sit in your stack.

Your systems keep running. Ours reads what they already produce, and hands back a ranked queue and a drafted case.

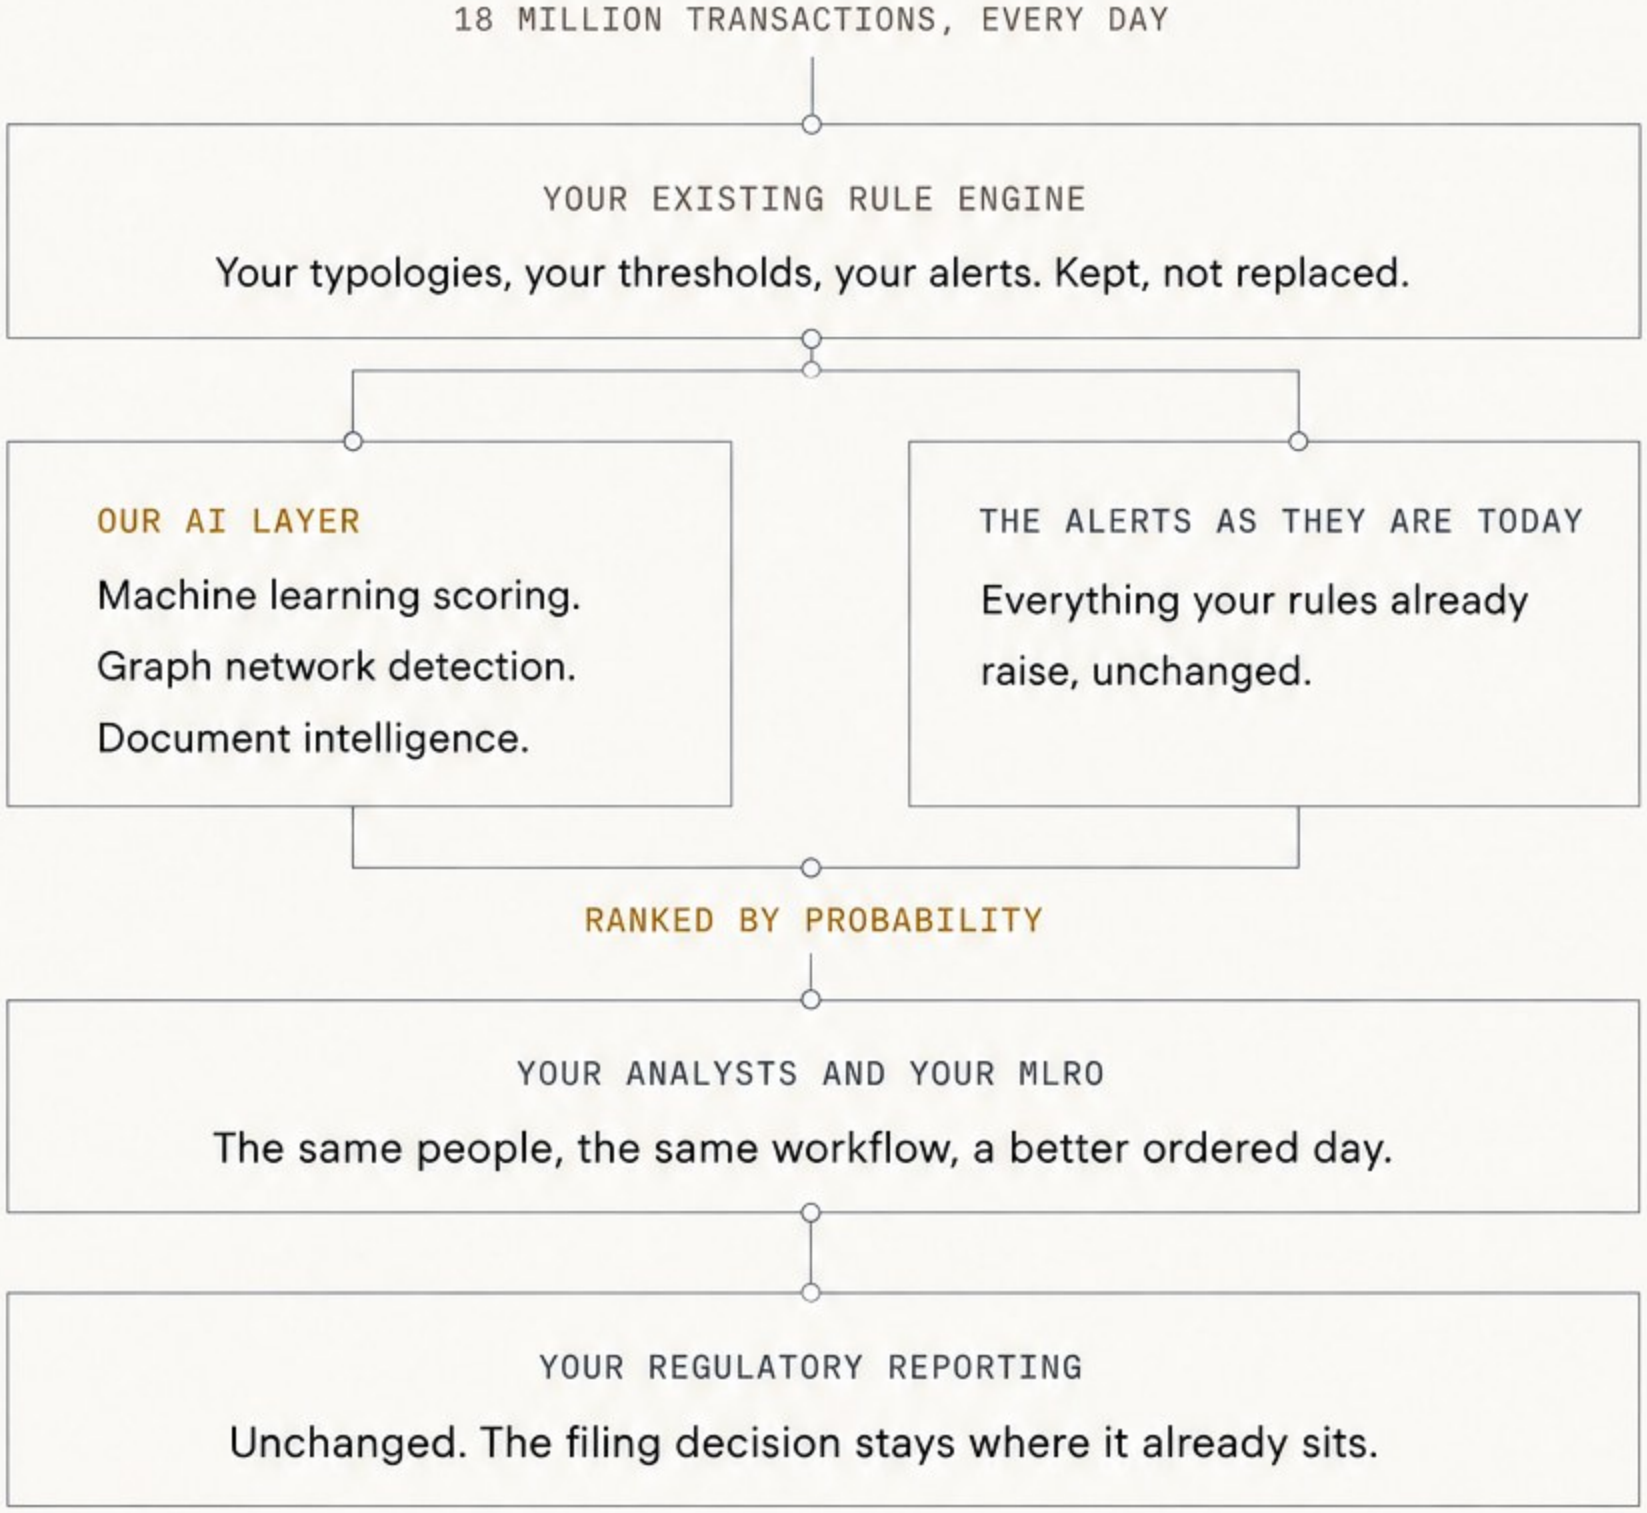

Every case your analysts confirm tunes both the model and your rules.

REDRAWN FROM A PRODUCTION DEPLOYMENT RUNNING ROUGHLY 18 MILLION TRANSACTIONS A DAY.

We deliver inside your perimeter.

Where the system runs, which model sits behind it and who holds the keys are all your decisions, and they stay your decisions.

SOVEREIGN PRIVATE CLOUD, OR ON PREMISES	The deployment decision stays with the bank at every stage, including the decision to change it later.
MODEL AGNOSTIC ARCHITECTURE	Commercial, open weight or partner models behind one interface. Swap the model without rebuilding the workflow around it.
DATA RESIDENCY AND KEY CONTROL	Transaction and supervisory data stays inside your organization, with the bank holding its own keys.
EXPLAINABILITY AND AUDIT BY DEFAULT	Not a feature you switch on. It is the design discipline behind every capability in this brochure.

HOW WE WORK

A long term partnership, delivered through a phased, low risk implementation that prioritizes business outcomes, regulatory compliance and capability your team keeps.

The deployment decision is yours at every stage, including the decision to change it.

Start with one. Grow into the engine.

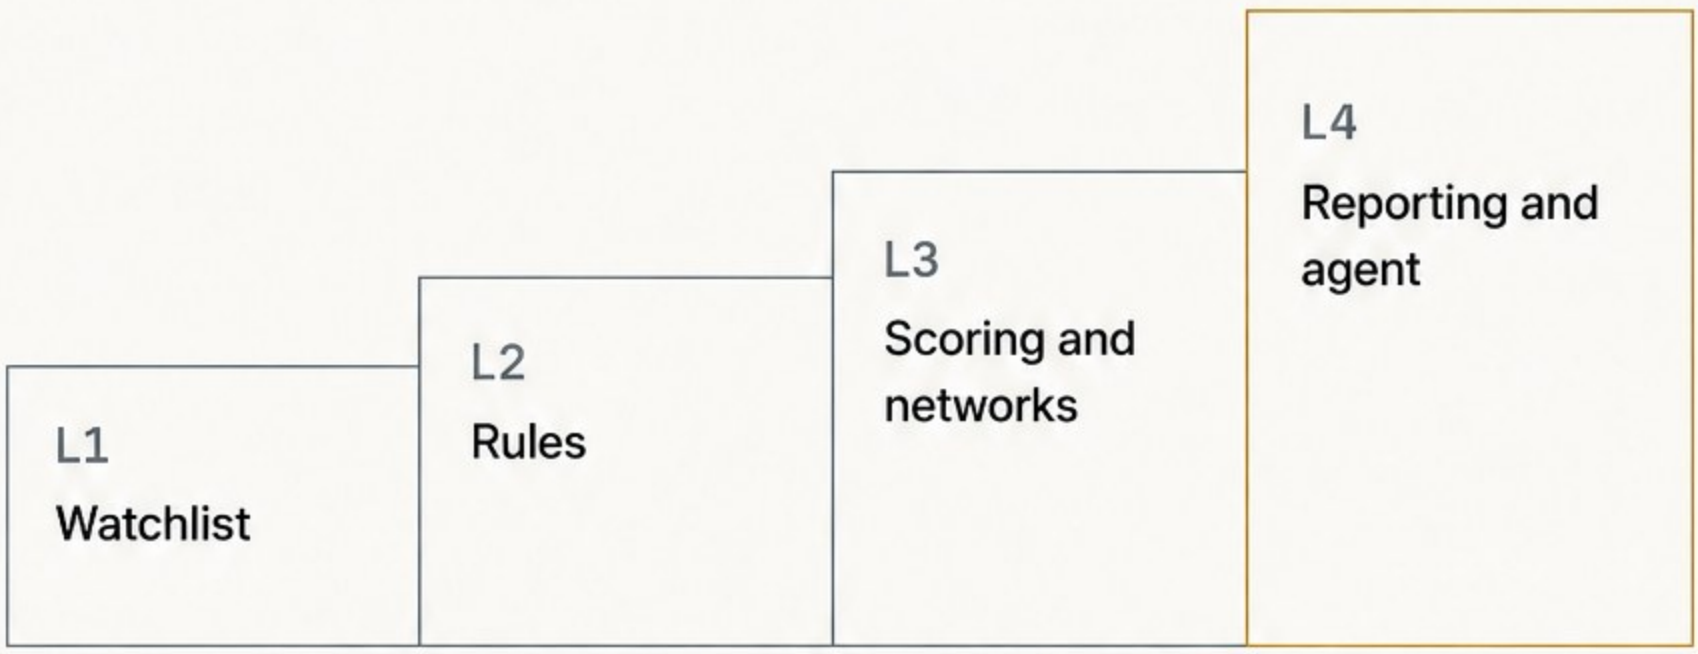

L1 WATCHLIST AND SANCTIONS	Static matching, adapted to your regulator. Under two weeks.
L2 THE RULE ENGINE PACKAGE	No code configuration for compliance officers.
L3 SCORING AND NETWORKS	ML case scoring, recall expansion, syndicate tracing.
L4 REPORTING AND THE AGENT	Independent of L1 to L3. Runs on infrastructure you own.

Each level stands on its own. Add one later and it joins the same engine, on the same audit trail.

Nothing you start with is thrown away.

Proven at scale.

Three engagements, each answering a different question a bank actually asks.

WILL IT SURVIVE IN MY BANK?

Core banking at a state owned commercial bank serving 192 million clients. 200 million transactions a day, 20,000 per second, 65 TB of logs daily with zero delay. An infrastructure engagement, not an AML one, and we show it deliberately.

DOES IT CATCH MORE THAN MY RULES?

AML in production at a top five retail bank: 200 million customers, $2.1 trillion in assets, 700 million transactions analyzed, trained on 330,000 confirmed cases. AUC 0.949 on held out validation.

DO I HAVE TO REPLACE ANYTHING?

A machine learning layer placed on top of an existing rule engine at a mid size joint stock commercial bank. Coverage up 33%, precision 26.8 times the rule layer alone, and 96% fewer false positives per confirmed case.

We build the floor a bank stands on, and we build the engine that runs on top of it. We either stay out of an industry, or we go deep.

What we do not claim.

WE DO	WE DO NOT
prioritize and improve alert investigation	replace your detection stack, and we will argue if you want us to
bring a real network platform and real patterns	arrive with a model calibrated to your market on day one
draft suspicious activity and transaction reports	approve or submit them. Your reviewers do. Every time
quote a document workflow against your numbers	quote anything before we have your volumes and handling times
scope trade based laundering where data exists	have trade data sources today

We would rather tell you where the claim stops than have you discover it in month four.

See these numbers on your own book, not ours.

Everything in this brochure happened somewhere else, to someone else's data. You should treat it accordingly. So the ask is not a signature.

A proof of concept, on your data.

SCOPE	Agreed on your own transaction data.
TEAM	A joint working team and your compliance stakeholders.
CONFIGURATION	Typologies, thresholds and workflows set to your bank.
A DATE	Set before anything begins.

We backtest against decisions your own people already made: cases they closed, cases they escalated, cases they filed, before anything touches production. You will know whether this works on your book before you have changed how it is run.

Your best reviewer's judgment is already correct. It is just not written down anywhere, and it does not scale to two in the morning.

Your Vision.
Your Stack.
Your Control.
Your Intelligence.

Our Knowhow Transfer.

AMANAH

SOVEREIGN AI INFRASTRUCTURE