

A Compliance Officer's Playbook

15 Real-World Compliance Problems Solved
with AI-Powered Analytics

by

H3M Analytics Inc.

Executive Summary

Financial-crime typologies mutate faster than any single rule or team can keep pace. To stay ahead, compliance programs now hinge on modular AI—tools that learn from data, surface hidden networks, and knit fragmented evidence into one coherent story. *A Compliance Officer's Playbook* distills those capabilities into fifteen field-tested use-cases drawn from banks, money-service businesses, and fintechs on four continents. Each vignette follows a common arc—problem, step-by-step solution, H3M module-mix, and quantified results—so readers can quickly translate lessons into their own controls.

Across the pages you'll see single-module wins (e.g., Sanctions Miner slashing false matches) and cross-module “fusion” workflows that expose complex trade-based laundering or mule rings. Together they illustrate a modern operating model: human judgment steered by machine-driven triage, network analytics, continuous active-learning feedback, and a 360° risk view. The matrix below summarizes which modules power which story, helping you jump straight to the examples most relevant to your current challenge.

Use-Case	Suspect Miner	Active Learner	Link Miner	Scenario Manager	Backlog Miner	Sanctions Miner	SWIFT Miner	OCR Miner	KYC Miner	KYC 360 Dashboards	RUMI (Adverse Media)	Federated Learner
Hidden-Risk ML												
Gambling Ring Network												
Optimizing AML Scenarios												
Real-Time Adaptation w/ Feedback												
Daily Automated Detection												
Clearing Alert Backlog												
Enhanced Sanctions Compliance												
AI-Powered SWIFT Screening												
TBML Detection via Documents												
Dynamic Customer Risk Scoring												
360° Holistic Risk Oversight												
Uncovering Hidden Risks with Adverse Media												
Community Defense via Federated Learning												
Multi-Module ML Stops Complex Scheme												
Multi-Module Crack-Down on TBML												

<u>USE CASE 1: USING MACHINE LEARNING TO UNCOVER HIDDEN RISKS</u>	<u>4</u>
<u>USE CASE 2: UNCOVERING AN ILLICIT GAMBLING RING.....</u>	<u>6</u>
<u>USE CASE 3: OPTIMIZING AML SCENARIOS WITH ACTIVE LEARNING</u>	<u>8</u>
<u>USE CASE 4: ADAPTING IN REAL-TIME WITH ACTIVE LEARNING FEEDBACK</u>	<u>10</u>
<u>USE CASE 5: AUTOMATING DAILY AML DETECTION WITH SCHEDULED MINER</u>	<u>12</u>
<u>USE CASE 6: CLEARING ALERT BACKLOG WITH BACKLOG MINER PRIORITIZATION .</u>	<u>14</u>
<u>USE CASE 7: ENHANCING SANCTIONS COMPLIANCE WITH AUTOMATED SANCTIONS MINER</u>	<u>16</u>
<u>USE CASE 8: SLASHING SWIFT SCREENING FALSE ALARMS WITH AI-POWERED SWIFT MINER</u>	<u>18</u>
<u>USE CASE 9: DETECTING TRADE-BASED MONEY LAUNDERING WITH OCR MINER....</u>	<u>20</u>
<u>USE CASE 10: REFINING CUSTOMER RISK SCORING WITH KYC MINER</u>	<u>22</u>
<u>USE CASE 11: HOLISTIC RISK OVERSIGHT WITH KYC 360° DASHBOARDS</u>	<u>24</u>
<u>USE CASE 12: UNCOVERING HIDDEN RISKS WITH RUMI ADVERSE MEDIA SEARCH</u>	<u>27</u>
<u>USE CASE 13: COMMUNITY APPROACH – FIGHTING FINANCIAL CRIME VIA FEDERATED LEARNING</u>	<u>29</u>
<u>USE CASE 14: MULTI-MODULE AI WORKFLOW STOPS COMPLEX MONEY LAUNDERING SCHEME.....</u>	<u>31</u>
<u>USE CASE 15: MULTI-MODULE CRACKDOWN ON TRADE-BASED MONEY LAUNDERING</u>	<u>33</u>

Use Case 1: Using Machine Learning to Uncover Hidden Risks

Problem Definition: MSB-1, a high-volume money transfer firm, was missing complex money laundering patterns with its rule-based system. False positives were overwhelming analysts, yet some sophisticated layering went undetected. They needed an AI-driven approach to identify **unusual transactional behaviors** indicative of illicit gambling and layering that rules couldn't catch.

Solution Steps:

1. **Data-Driven Model Training:** The team loaded historical transactions and alert outcomes into KROTON **Suspect Miner**. They created a new machine learning “miner,” engineering 100+ risk features from CRM (customer profiles) and transaction data.
2. **ML Anomaly Detection:** Suspect Miner’s adaptive algorithms scoured the data for hidden patterns. It identified a cluster of accounts with **unusual funds flows** (rapid in-and-out transfers just under reporting thresholds) that had not triggered any rule.
3. **Analyst Review & Refinement:** Compliance analysts reviewed the flagged accounts, confirming many as true risks. They provided feedback in the system, which Suspect Miner incorporated to fine-tune the model. The integrated Active Learning mechanism ensured the model continuously improved with each analyst input.
4. **Deployment & Monitoring:** The trained model was deployed on live data. Suspect Miner now scores transactions daily, flagging high-risk entities. Analysts receive a prioritized list of suspects, each with an **ML risk score** and explanation of contributing factors.

Tools Used:

- **KROTON Suspect Miner:** Machine-learning module that analyzes transaction data to detect **anomalous patterns beyond defined scenarios**. It allows customizable feature engineering and model tuning to adapt to MSB-1’s specific risk profile (e.g., spikes in send/receive frequency).
- **Active Learning Engine:** The platform’s ability to **continuously retrain** on new feedback ensures emerging typologies are quickly learned, keeping detection capabilities up to date.

Kroton Suspect Miner - Miner Training tolga (Master, Admin, Analyst)

SELECT MINER

tt NEW SELECT DELETE

SELECT ALERT DATA

☒ Filter By rules

Select All

☐ SOMERULE

☐ EXCEL

Start Date: 2/1/2018 End Date: 2/8/2023

White List: Leaving Clear

☒ Person ☐ Company ☐ Personal Company ☐ Joint Account

SETTINGS

Status:

KRT_Clear	<input type="radio"/> SAR	<input type="radio"/> Clear	<input checked="" type="radio"/> Unknown
RED	<input type="radio"/> SAR	<input type="radio"/> Clear	<input checked="" type="radio"/> Unknown
2	<input type="radio"/> SAR	<input checked="" type="radio"/> Clear	<input type="radio"/> Unknown
KRT_Suspect	<input type="radio"/> SAR	<input type="radio"/> Clear	<input checked="" type="radio"/> Unknown
KRT_No-SAR	<input type="radio"/> SAR	<input type="radio"/> Clear	<input checked="" type="radio"/> Unknown
TAMAMLANDI	<input type="radio"/> SAR	<input type="radio"/> Clear	<input checked="" type="radio"/> Unknown
6	<input checked="" type="radio"/> SAR	<input type="radio"/> Clear	<input type="radio"/> Unknown

Relative cost of (PSN/FP): 10

Alert Date Fields: ☐ Alert date ☒ Decision date

☒ Enable Random Forest Parameters

☐ Optimize RF parameters

Parameter	Value
Number of estimators	100
Maximum depth	99
Minimum # of samples to split	3
Minimum # of samples to leaf	4

ADD CRM FEATURE ADD TRX FEATURE

☐ ESTABLISHMENTDATE EDIT

☐ CREATEDAT REMOVE

☐ INCOME PREVIEW

☐ NACE

☐ NATIONALITY

☐ Transaction count 1d to 1w

☐ Outgoing total 1m to 6m

Outcome Achieved:

- **High Precision Alerts:** The Suspect Miner achieved an **~98% true positive alert rate**, dramatically enhancing detection accuracy. Alerts are now far more meaningful, with up to a 90% reduction in false positives, allowing the team to focus on genuine threats.
- **New Typologies Detected:** MSB-1 uncovered previously missed schemes (e.g. an illegal gambling ring moving funds in small increments) that legacy rules never flagged. These insights led to the filing of multiple SARs that might not have been caught otherwise.
- **Operational Efficiency:** By slashing noise, analyst workload dropped significantly. The compliance team can handle alerts with existing staff, spending time on in-depth investigations instead of clearing benign alerts.
- **Dynamic Compliance:** The system continuously adapts to evolving criminal techniques. As patterns shift, the model self-adjusts with each feedback loop, ensuring **robust, up-to-date compliance controls**.

Use Case 2: Uncovering an Illicit Gambling Ring

Problem Definition: Bank-A noticed hints of an **illegal gambling ring** – numerous small accounts pooling money to a few recipients – but individual transactions appeared innocuous. Traditional monitoring couldn’t connect these dots, as **no single account’s activity looked suspicious in isolation**. The challenge was to reveal the hidden network of relationships behind these transactions.

Solution Steps:


1. **Seed Identification:** Investigators input a set of initially suspicious accounts (e.g. a handful of frequent small senders and a central receiver) into **KROTON Link Miner**. They defined a date range covering the last 6 months of transactions to focus the analysis.
2. **Network Forming & Filtering:** Link Miner pulled all transactional links among those accounts and their counterparts. The team applied *pre-filters* to ignore very low-value transactions and one-off links and used the **“strong connections only”** mode to focus on significant, repeated flows. They also filtered by a rule set (omitting clearly legitimate salary payments, for example) to home in on suspicious activity.
3. **Running Link Analysis:** The Link Miner algorithm ran on the filtered data, mapping out the network graph. It **ranked each node (customer) by suspicious connectivity**, identifying key hubs where many funds converged. Within minutes, Bank-A saw a web of ~20 individuals all funneling money to a single organization masquerading as a “sports club.”
4. **Visualization & Investigation:** The tool plotted an interactive network graph. Analysts examined this visualization, spotting clusters – e.g., several bettors (nodes) all connected to the “club” account (central node). They drilled down on connections and found some nodes were linked via shared phone numbers and referees, corroborating the ring hypothesis.
5. **External Data Overlay:** To enrich the picture, the team overlaid external info. Link Miner integrated data from public leaks – it revealed one connected entity appeared in the Panama Papers as owner of an offshore betting site. This reinforced that the network was part of a larger illicit operation.

Tools Used:

- **KROTON Link Miner:** Network analysis module that finds **suspicious relations among entities**. It constructed a detailed map of Bank-A’s customer transaction network, uncovering relationships that were not apparent on a per-account view.
 - *Pre- & Post-Filters:* Allowed analysts to filter input by transaction traits (e.g. only include accounts triggering certain rules) and output by network metrics, honing in on meaningful links.
 - *High-Speed Graph Engine:* Scaled to analyze millions of nodes and links in under an hour, meaning Bank-A’s entire customer network could be processed rapidly.
 - *“Strong Connections” Mode:* Focused on significant recurring flows, effectively highlighting the **“follow the money” trail** while pruning noise.
 - *Visualization & Third-Party Data Integration:* Provided clear network graphs and integrated external datasets (e.g., known shell companies from leaks) for broader context.

Outcome Achieved:

- **Hidden Network Exposed:** Bank-A definitively uncovered an illicit gambling ring. Over a dozen gambler accounts and the central “club” account were linked in one view. This **complex network, which eluded rule-based detection, was exposed by Link Miner’s AI-driven graph analysis.**
- **Swift Action:** What could have taken months of manual tracing was achieved in a single analysis session. The system processed the network in ~45 minutes and presented investigators with clear evidence, **significantly reducing investigation time.** Bank-A quickly froze the key accounts and filed a comprehensive SAR, likely preventing further illicit transactions.
- **Efficiency & Focus:** The network visualization guided analysts straight to the “**central nodes**” (top suspicious actors), ensuring they prioritized the true orchestrators of the scheme. This targeted approach improved the efficiency of the investigation team and avoided wasting effort on peripheral transactions.
- **Strengthened Controls:** Bank-A implemented new monitoring measures as a result – for example, enhanced scrutiny on groups of accounts interlinked through common beneficiaries. The success of Link Miner here set a precedent to routinely use network analysis for any **suspected rings or collusion**, greatly boosting the bank’s ability to tackle organized financial crime.



Use Case 3: Optimizing AML Scenarios with Active Learning

Problem Definition

Bank-B's legacy rule set began generating a flood of nighttime ATM-deposit alerts: **nine out of ten** were cleared by investigators. Analysts spent hours triaging taxi-driver wage deposits while genuine money-laundering risks hid in the noise. The compliance team needed a smarter way to refine its scenario—without losing true positives.

Solution Steps

1. **Scenario Build (Version 1)**
 - **Rule** – “Frequent Night-Time ATM Cash Deposits”
 - Trigger: ≥ 10 cash deposits between 22:00-05:00 within any 7-day window.
 - Initial score weights: deposit count 60 %, night-time factor 40 %.
2. **Active-Learning Iteration #1**
 - **Active Learner** surfaced “cleared-but-predicted-suspicious” cases.
 - Analysts discovered most were *taxi drivers* ending shifts and depositing fares for security purposes.
 - **Scenario Manager tweak:** occupation = “Taxi driver” \rightarrow **-20 score** (negative weight).
3. **Active-Learning Iteration #2**
 - Some taxi drivers later showed up in FIU feedback as SARs.
 - Deeper review revealed drug dealers coercing taxi drivers to act as money mules: cash in at night, funds wired out within 24 h.
 - **Scenario Manager tweak:** *IF ≥ 95 % of deposits leave the account < 24 h \rightarrow +20 score* (overrides taxi discount).
4. **A/B Validation**
 - Scenario Manager ran the new logic in parallel with the legacy rule set over four weeks to compare alert quality, workload, and missed SARs.

Tools Used

- **Scenario Manager** – versioning, weighting and A/B-testing of AML rules.
- **Active Learner** – generates four classes of questions (“cleared-but-predicted-suspicious”, etc.) so analysts’ feedback continuously reshapes scenario logic.

SELECT SCENARIO

☐ AI Risk Score
☒ High Frequency ATM Deposits
☐ Federated Learning Test_1
☐ Sanctions Evasion 1

NEW

SELECT

DELETE

THRESHOLD

99

STATUS

Active

TYPOLGY 1

Select Typology1

SK-01-Tefecilik

TYPOLGY 2

Select Typology2

SK-02-POS Tefeciligi

DEFINITION

Definition

göce yasma

CREATED DATE

2025-02-02

MODIFIED DATE

2025-04-09

MODIFIED BY

urku

ADD CRM FEATURE

	Feature Name	Operator	Feature Details	Pre-filter	Score	Action
TYPE	TYPE	IN	PERSON	<input checked="" type="checkbox"/>		
NACE	Night ATM Deposit Count	>	10	<input type="checkbox"/>	100	
	NACE	IN	Taksicilik	<input type="checkbox"/>	-20	
	Ratio of outgoing incoming	>	0.95	<input type="checkbox"/>	20	

SAVE

SAVE AS

ADD TRX FEATURE

Night ATM Deposit Count

Ratio of outgoing incoming

Outcomes Achieved

- **False-Positive Rate:** ↓ 78 % (taxi-only alerts largely eliminated).
- **True-Positive Rate:** ↑ from 55 % to 92 % (money-mule pattern now detected).
- **Analyst Efficiency:** ≈ 120 hours/month released for deeper investigations.
- **Regulatory Assurance:** Clear audit trail of three scenario versions, each backed by Active-Learning evidence and measurable KPI gains.
- **Data-Driven Culture:** Success prompted quarterly scenario reviews using the same Active Learning loop, embedding continuous improvement into Meridian's AML program.

Use Case 4: Adapting in Real-Time with Active Learning Feedback

Problem Definition: Bank-C noticed that even with advanced models, **new fraud patterns and edge cases** were emerging that weren't correctly classified. Traditional retraining cycles were too slow. They needed a way for their AML system to *learn continuously from analyst expertise* – catching unknown typologies and reducing misclassifications on borderline alerts. In short, the bank sought to close the loop between human investigators and the AI model in real time.

Solution Steps:

1. **Active Learning Test Creation:** The compliance team initiated an Active Learning cycle using H3M's **Active Learner** module. An **Active Learning Test** was generated, where the system pulled together a set of tricky cases for analysts to label. These included:
 - *False negatives:* past transactions that analysts had manually found suspicious but the model predicted as “clear”.
 - *False positives:* alerts the model flagged as suspicious but analysts consistently cleared.
 - *Uncertain cases:* transactions where the model was ~50/50 unsure (near the decision boundary).
 - *Learning-maximization cases:* a few unlabeled instances that the algorithm believed would most improve the model if clarified.
2. **Question Deployment to Analysts:** The module assembled these cases (e.g., 20 customer transaction profiles) and assigned them to senior compliance analysts via the Active Learning Test interface. Analysts received notifications and accessed the test in the system.
3. **Analyst Labeling:** For each case, analysts provided a definitive label – “**SAR**”, “**Clear**”, or “**Anomaly**” (meaning something about the case invalidated prior assumptions) – along with comments explaining their reasoning. For example, an analyst might label a borderline case as “Clear” noting it was a known legitimate remittance pattern, or mark “SAR” for an odd transaction citing a new typology.
4. **Model Update:** Once all questions were answered and the test completed, Active Learner fed these labels back into Suspect Miner. The ML model retrained incorporating the newly labelled data. Misclassified examples now corrected the model's internal decision boundary (e.g., recognizing that pattern X is actually low risk, while pattern Y is indeed suspicious).
5. **Iterative Improvement:** The updated model was deployed, immediately adjusting alert outputs. The system scheduled such Active Learning tests monthly. Over successive rounds, Bank-C's model became significantly more accurate, having effectively **crowd-sourced expert knowledge into the AI**.

Tools Used:

- **Active Learner Module:** A proprietary H3M module enabling **human-in-the-loop machine learning**. It generates targeted questions on four case types (false positive, false negative, uncertain, high-impact new cases) for analyst review.

- **Suspect Miner Integration:** The Active Learner seamlessly integrates with the Suspect Miner model. Once analysts complete the test, results are automatically applied to retrain the model on those answers, **improving prediction accuracy on the next run.**

[illegible]

Outcome Achieved:

- **Enhanced Detection Accuracy:** Bank-C saw measurable gains – after a few Active Learning cycles, the true positive rate of its alerts improved substantially while false positives dropped. The model learned to catch new schemes (some typologies were completely unknown before) and to stop flagging certain innocuous patterns. This aligns with H3M’s observed effect: using active learning **drastically boosts detection of previously unknown patterns**.
- **Continuous Adaptation:** The AML program moved from static annual model updates to a **continuous improvement process**. With expert feedback regularly ingested, the system stays in sync with the latest criminal tactics. For instance, when COVID-era scam patterns emerged, analysts labeled a few examples and the model quickly adapted within weeks, not years.
- **Analyst–AI Synergy:** Investigators felt empowered – their expertise wasn’t just resolving one alert at a time, but also making the system smarter with each input. This led to a virtuous cycle: the better the model got, the fewer trivial alerts analysts saw, freeing them to focus on complex cases and give quality feedback, further improving the model.
- **Reduced Operational Burden:** Over time, Bank-C dramatically reduced the manual effort needed for alert review. With the model doing a better job upfront, only truly ambiguous cases reach human eyes. The **workload reduction and efficiency gains** were significant, on the order of a 50% decrease in daily alerts to review. This outcome mirrors industry results where active learning **decreases manual effort and increases detection of new threats**.
- **Robust & Up-to-Date Compliance:** Perhaps most importantly, Bank-C’s AML controls became more **robust against evolving risks**. The dynamic learning approach ensured the bank could catch emerging typologies (e.g., novel fraud rings, new money laundering channels) that would have slipped past a static system. This proactive stance significantly boosts confidence from regulators that Bank-C can handle ever-changing financial crime risks.

Use Case 5: Automating Daily AML Detection with Scheduled Miner

Problem Definition: MSB-2, a fast-growing e-money wallet provider, needed to monitor transactions continuously. Previously, they ran AML checks in batches infrequently, causing delays in flagging suspicious activity. The compliance team sought an **automated scheduling solution** to ensure their machine learning models (Suspect Miner) run on every new day's data without manual intervention, providing timely alerts and consistent coverage.

Solution Steps:

1. **Schedule Configuration:** MSB-2's analysts used the **Scheduled Miner** module to set up daily runs of their trained Suspect Miner model. They selected the specific miner model (tuned for MSB-2's risk patterns) and configured it to execute every night at 2 AM after that day's transactions were in.
2. **Parameter Settings:** In the scheduler settings, they defined run parameters: a maximum of 50 alerts per run (to control output volume) and a data "lag" of 1 day (meaning each run analyzes transactions up to the end of the previous day). They also utilized the *Worker/Learner ratio* setting – e.g., 80% normal detection vs 20% active learning cases – so each run would include a few Active Learning alerts for continuous model refinement.
3. **Activation:** Once saved, the schedule was active. The Suspect Miner model now **automatically runs each day** with the specified settings. No manual trigger is needed – the system kicks off the job at 2 AM, processes all new transaction data, and applies the ML model to score customers.
4. **Automated Alerts & Listing:** After each run, the results (suspicious customers with scores) are available first thing in the morning in the **Suspect Listing** interface. Compliance officers come in and immediately see, for example, "85 new suspects scored overnight," each with a risk score between 0.83 and 1. They filter these results by date and model ID to view the latest batch and begin investigations right away.
5. **Ongoing Management:** The team monitors the Scheduled Miner via the Task Manager, which shows status (completed, running, etc.) of each scheduled run. They can make adjustments if needed – for instance, temporarily increase the alert cap if expecting a spike in activity, or pause the schedule during a system maintenance window.

Tools Used:

- **Scheduled Miner:** Allows **automated, periodic execution** of Suspect Miner models. MSB-2 set it to run daily, but it supports flexible intervals (weekly, hourly, etc.). Key features include:
 - Setting maximum alerts per run to manage review workload.
 - Specifying data range or lag, ensuring the model always runs on the freshest data (e.g., "up to yesterday's transactions") for timely detection.
 - Option to blend in Active Learning cases each run (via worker/learner ratio) so the model continues learning even during scheduled execution.

- **Suspect Listing & Notification:** Once a scheduled run completes, alerts are populated in the suspect listing window automatically. MSB-2's team also set up email notifications on run completion, so they're alerted if a run produces high-severity hits.
- **Task/Activity Log:** The platform logs each scheduled run, which is useful for audit. If a run fails or is delayed, the Task Manager highlights it, enabling quick troubleshooting. (In practice, the Scheduled Miner proved reliable, running every night without incident.)

Name	Status	Last Run
Run Agent Miner	Failed	01-01-2025 00:00
Sanction Matcher	Active	01-01-2025 00:00
Sanction Matcher	Active	01-01-2025 00:00

DELETE PAUSE FORCE ACTIVATE NEW EDIT

EDIT SCHEDULE

Select Miner:
☐ SELECTED
☐ SELECTED FOR REVIEW
☐ SELECTED (Source: Select Miner)
☐ SELECTED (Source: Select Miner)
☐ SELECTED (Source: Select Miner)
☐ SELECTED (Source: Select Miner)

Every: Day

Select White List: ▼

Max Alerts: 85 alerts/day

Data Lag: 1 days

Continue If: ☐

Runners: 100 - 500

SAVE

Outcome Achieved:

- **Real-Time Vigilance:** MSB-2 now **flags suspicious activity within hours**, not weeks. Every day's transactions are analyzed overnight, so potentially illicit activity is caught almost immediately the next morning. This dramatically shortens the window in which bad actors can operate undetected, improving the prevention and interdiction of laundering attempts.
- **Consistent Monitoring:** The automated schedule ensures no days are missed – compliance monitoring is 24/7. Previously, if a staff member forgot to run a report or took vacation, gaps could occur. Now MSB-2's AML is on autopilot, providing **uninterrupted coverage** of transactions.
- **Reduced Manual Work:** The compliance team saved considerable manual effort. What used to be a tedious daily task (running queries, generating alert files) is fully automated. This reduction in manual workload allows analysts to **focus on higher-value activities** like investigating alerts and refining scenarios, rather than data wrangling.
- **Timely Investigations:** Alerts are fresher and more relevant. Investigators can act on them when the transactions are recent (often same-day), which makes inquiries with customers or counterparties more effective. MSB-2 noted that responding quickly often led to better cooperation (e.g., freezing funds before they left the system).
- **Scalability:** As MSB-2's volume grows, the Scheduled Miner setup easily scales. They've configured it to handle increasing transactions without additional headcount. This has **enhanced efficiency** and given management confidence that compliance can keep up with business expansion. In essence, MSB-2 achieved a **reduced workload and increased efficiency** in transaction monitoring through automation, while maintaining (even improving) detection performance.

Use Case 6: Clearing Alert Backlog with Backlog Miner Prioritization

Problem Definition: Bank-D had accumulated a **huge backlog of unreviewed AML alerts** from its old rule-based system – thousands of alerts were generated faster than the team could investigate. Many were likely false positives but hidden among them could be true suspicious cases. Each alert was a customer that had triggered a rule and was marked “unknown” (not yet dispositioned). The bank needed a way to triage this backlog: identify which alerts to tackle first (and which might be safely deprioritized) to reduce risk and workload.

Solution Steps:

1. **Miner Selection:** The bank took a trained and approved Suspect Miner model (one that had learned from Legacy’s past data) and applied H3M’s **Backlog Miner**. They selected this model in the Backlog Miner interface and specified the date range covering all outstanding alerts from the last year.
2. **Scope Definition:** Analysts narrowed the scope by filtering to scenarios of interest. For example, they focused on backlog alerts from two rules historically prone to false positives (like a threshold rule on cash deposits), to rank those first. This option to filter by scenario meant they could tackle the backlog in logical chunks (prioritize the rules likely hiding true positives).
3. **Ranking Execution:** They clicked “Rank Backlog.” The Backlog Miner algorithm processed each alert (each customer in the backlog) with the ML model and assigned a **suspiciousness score from 0 to 1** to every pending alert. A score of 1.0 indicates the alert is very similar to known suspicious cases, while 0 means it looks innocuous. This ranking, along with customer IDs and alert info, was output to a CSV report for easy review.
4. **Review of High-Scorers:** The team sorted the backlog by the score. Immediately, patterns emerged – e.g., about 50 alerts scored above 0.8 (likely true positives). Analysts began investigating those high-scoring alerts right away. Many turned out to indeed warrant SARs (the model effectively prioritized alerts that the old system had flagged for good reason).
5. **Deprioritization of Low-Scorers:** Conversely, a large portion of alerts scored very low (near 0). Alerts below a certain threshold (e.g., <0.2) were deemed very likely false positives. The bank decided to **bulk close or deemphasize** those after a spot-check confirmed they were benign (e.g., repetitive known customer behavior). This freed the team from having to manually investigate every single alert in the backlog.
6. **Iterative Clearance:** Using this prioritization, Bank-D systematically worked through the backlog: investigate the top tier thoroughly (which contained almost all the true risks) and safely dismiss the bottom tier. For middle-range scores, they applied normal investigative judgment. Over a few weeks, the once-daunting backlog was essentially cleared or re-risk-rated.

Tools Used:

- **Backlog Miner:** This module leverages an existing ML model to **re-score and rank backlog alerts** (alerts generated by third-party scenario-based systems). Key capabilities used by Bank-D:

- | # | DISPLAY | AUDET2 | CUSTOMER | AGE (CREATED) | TRF_OFF_REM_Age_Lambda | TRF_OFF_REM_Age_Lambda | TRF_Lat_Month_Absolute | NATIONALITY | TR_GENRES | TR_GENRES | TR_Probability of SA | Type 1 | Type 2 |
|----|--------------|--------------|-------------|---------------|------------------------|------------------------|------------------------|-------------|-----------|-----------|----------------------|-----------------|------------------|
| 1 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 2 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 3 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 4 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 5 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 6 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 7 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 8 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 9 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 10 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 11 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 12 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 13 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 14 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 15 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 16 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 17 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 18 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 19 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 20 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 21 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 22 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 23 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 24 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 25 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 26 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 27 | 18709712100A | 18709712100A | 18709712100 | 64 | 18 | 0 | 0 | 0 | 0 | 0 | 0.99942130 | 18-22 Age Group | 18-21 Sex Female |
| 28 | 18709712100A | 18709712100A | | | | | | | | | | | |

Use Case 7: Enhancing Sanctions Compliance with Automated Sanctions Miner

Problem Definition: Bank-E deals with thousands of customer onboarding and payment transactions daily that must be screened against sanctions lists (OFAC, UN, EU, etc.) and PEP lists. Maintaining up-to-date coverage was challenging – lists update frequently, and **false positive name matches** were overwhelming (common names like "Mohammed Ali" triggered constant alerts). The bank needed a **comprehensive, accurate sanctions screening solution** that could keep pace with updates and drastically reduce false hits.

Solution Steps:

1. **Deployment of Sanctions Miner:** The bank implemented H3M's **Sanctions Miner** and integrated it with their core systems. They configured it to pull in all major global and regional sanctions lists (OFAC, EU, UK-HMT, UN, plus local watchlists) with **daily automatic updates**. This ensured their screening database was always current with evolving regulations.
2. **Integration & Modes:** The solution was set up in multiple modes:
 - **Real-time API Screening:** Every new customer or payee added was screened instantly via the Sanctions Miner API. If a match (or near match) was found, the onboarding was flagged for review.
 - **Batch Delta Scans:** Each night, the system did a delta scan – screening any changes (new transactions, account updates) against the updated lists. This covers scenarios like an existing customer becoming sanctioned overnight.
 - **Full Periodic Scans:** A full scan of the entire customer base was scheduled monthly as a safety net, though delta scans captured interim changes.
3. **Advanced Name Matching:** The bank enabled Sanctions Miner's **NLP-based risk scoring** for name matches. Instead of simple exact or phonetic matches, the system uses natural language processing to evaluate context and similarity. For example, it can differentiate between "John Smith" the sanctioned narcotics trafficker and other John Smiths by analyzing additional data (address, DOB, etc.) and then assigning a risk score to the match.
4. **Customization:** Compliance configured custom whitelist entries for known innocuous matches (e.g., internal employees or common false positive names). They also added a few custom local blacklist entries for persons of interest (not on official lists yet). The system seamlessly incorporated these definitions, preventing repeat false alerts and covering institution-specific risks.
5. **Screening in Practice:** When transactions were processed, e.g., a wire transfer, Sanctions Miner screened all parties (originator, beneficiary). If a name hit occurred above a certain risk score threshold, it created an alert in the case management system. One example: "Robert Mugabe" as a beneficiary triggered an instant high-severity alert (OFAC hit), whereas "Roberto Mugabi" triggered a low-risk alert that was auto-cleared by the system's scoring (NLP determined it was not the same person).
6. **Analyst Review Workflow:** For any flagged hits, analysts used the Sanctions Miner interface to see why (which list, what similarity score). With far fewer alerts, they could promptly investigate true matches. Integration with providers like Dow Jones also allowed them to pull detailed profile data on the matched entity with one click.

Tools Used:

- **Sanctions Miner:** End-to-end sanctions and watchlist screening solution. Key features leveraged:
 - **Comprehensive List Coverage:** Automatically aggregated all relevant sanctions and PEP lists (global + local) and updated them daily, eliminating the risk of missing a newly listed individual.
 - **Automated Updates & Integration:** Direct integration with data providers (Refinitiv, Dow Jones) for seamless list updates, and easy plug-in via API to Bank-E's onboarding and payments systems.
 - **Versatile Screening Modes:** Provided full scans, daily delta scans, and on-demand checks in real-time, giving the bank flexibility in how and when to screen.
 - **Advanced Matching & Scoring:** Used proprietary NLP algorithms to calculate match confidence scores. This drastically cut false positives by ignoring low-risk name coincidences and highlighting likely matches.
 - **Custom Lists Management:** Interface to maintain internal blacklists/whitelists. Bank-E added PEP relatives as high-risk custom entries and whitelisted certain repeated false matches, fine-tuning the process.

Outcome Achieved:

- **Low False Positives:** The intelligent scoring reduced the avalanche of false alerts. For instance, the bank saw an **80%+ reduction in sanctions screening false positives** in the first quarter. Alerts for common names dropped dramatically; only truly suspicious name matches now require attention. This aligns with reported outcomes where AI screening **"drastically reduced false positives, ensuring more precise and reliable sanctions screening"**.
- **Improved Detection & Compliance:** No sanctioned individuals slipped through. With daily list updates and thorough coverage, Bank-E remained continuously compliant with the latest sanctions. They even caught a few tricky cases – e.g., a client of a different name but identified as a **sanctioned entity's alias** was flagged correctly by the NLP engine (something the old system missed). The bank avoided potential violations and fines thanks to this rigor.
- **Efficiency Gains:** The first-line sanctions screening team was effectively **streamlined**. The bank was able to reassign or reduce headcount that previously manually cleared false hits. In one department, 5 out of 7 analysts were moved to other risk tasks as the AI took over initial review. This matches patterns where AI **replaces large first-line teams and lowers operational costs**.
- **Fast Throughput:** Screening became virtually instant. Onboarding customers went from a manual check taking several minutes to an automated process taking seconds, without sacrificing thoroughness. International payments processing speed improved since fewer transactions got held for manual OFAC review.
- **Regulatory Confidence:** During the next audit, Bank-E demonstrated their Sanctions Miner logs showing **all major lists monitored daily** and the dramatic false positive drop. Regulators were impressed with the proactive approach (e.g., using risk scoring, custom lists) and noted it as a best practice in meeting global sanctions compliance standards. The bank essentially turned sanctions screening into a strength, with rigorous coverage and efficient operations.

Use Case 8: Slashing SWIFT Screening False Alarms with AI-Powered SWIFT Miner

Problem Definition: Bank-F's international payments unit was flooded with sanctions screening alerts on SWIFT messages. The nature of SWIFT free-text fields (filled with names, references, etc.) led to **99.9%+ false positive rates** traditionally – almost every message with certain common terms triggered an alert. A team of 10 analysts manually reviewed these alerts, yet genuine sanction hits were extremely rare. This process was costly, slow, and prone to human error due to alert fatigue. Bank-F needed an advanced solution to **intelligently filter SWIFT messages**, reducing false hits while ensuring no true sanctions violations slip by.

Solution Steps:

1. **Introduction of SWIFT Miner:** Bank-F implemented KROTON's **SWIFT Miner**, a deep learning-based screening tool specifically for SWIFT payment messages. It was integrated into the SWIFT processing stream so that every incoming and outgoing message would pass through SWIFT Miner before reaching analysts.
2. **Model Training:** The SWIFT Miner's AI model was trained on a historical dataset of Bank-F's SWIFT messages, including which had been false positives and any true hits. Using this, the model learned to distinguish innocuous mentions from real risk. For example, it learned patterns like "CUBA STREET" (an address) vs. "Cuba – Ministry of Defense" (a likely sanction hit context).
3. **Real-Time Screening:** As messages flow, SWIFT Miner evaluates the text using **advanced deep learning NLP** algorithms. Instead of flagging based on simple keyword presence, it assesses context. A message mentioning "Sudan" in an unrelated context might be given a low risk score, whereas one referencing a known SDN (Specially Designated National) gets a high score. Only messages that exceed a risk threshold trigger alerts.
4. **Pilot Comparison:** Bank-F ran a 3-month pilot, letting SWIFT Miner screen in parallel with the human team. The results were telling: the AI consistently cleared benign messages that humans would have and flagged the same true hits the humans did. In fact, SWIFT Miner caught a subtle sanction-relevant message that the busy manual team initially overlooked. Metrics showed false positives were cut by over 90%, and the AI's decisions aligned with expert judgement almost perfectly.
5. **Operational Rollout:** After proving its effectiveness, Bank-F fully replaced the first-line manual screening with SWIFT Miner. Now, only a **tiny fraction of messages** (the ones SWIFT Miner flags with high risk) are routed to compliance analysts for review. All others pass through if cleared by the model, greatly accelerating throughput. Analysts were retrained to focus on the few complex cases rather than volume processing.

Tools Used:

- **SWIFT Miner:** A specialized AI module applying **deep learning to SWIFT message screening**. Key features utilized:
 - **Deep NLP Model:** Able to interpret the unstructured SWIFT message text and context, far surpassing traditional rule or keyword matching. It accounts for

misspellings, variations, and context around names to decide if a message likely concerns a sanctions target.

- **False Positive Reduction:** By design, it targets the extreme false positive problem (>99% FP). Bank-F saw the tool **drastically reduce false positives** in line with expectations, using semantic understanding to ignore irrelevant matches.
- **Integration Hooks:** Fully compatible with the SWIFT network and message formats. It slotted into Bank-F's payment workflow with minimal tech changes, intercepting messages via API and returning a clearance decision or alert.
- **First-Line Automation:** Proven capability to **replace a manual Level 1 screening team**, making it ideal for Bank-F's goal of redeploying human resources.
- **Monitoring Dashboard:** Provided oversight on SWIFT Miner's performance – e.g., daily counts of messages screened vs. alerted, and an interface to review the AI's rationale on flagged messages (showing which parts of text triggered risk). This gave compliance comfort and transparency on the AI's decisions.

Outcome Achieved:

- **Massive False Alert Drop:** Bank-F experienced an **over 99% reduction in SWIFT screening false positives**. Literally thousands of meaningless alerts per day were eliminated. Now, perhaps only 5–10 truly dubious messages a day require review instead of 1000+. The SWIFT Miner “set a new standard in reducing false positives and enhancing detection accuracy in international transfers”.
- **Cost & Efficiency Gains:** The 10-member team dedicated to scanning SWIFT alerts was **completely re-purposed or reduced**. This saved significant operational costs. Turnaround times improved too – previously each wire might be delayed for manual OFAC check, but now straight-through processing rates increased, benefitting customers and internal SLAs.
- **High Detection Confidence:** No loss in screening quality occurred. On the contrary, the deep learning model caught every true sanctions hit (and even potential ones humans might miss). Bank-F noted that during the pilot, **the AI matched or exceeded human accuracy** – it never missed a real hit that the humans caught, and even identified edge cases. This gave compliance leadership confidence to trust the AI fully.
- **Analyst Focus on Real Issues:** The few alerts that do come through are much more likely to be true problems. Analysts can now dedicate ample time to investigate those thoroughly (rather than speed-reading endless false positives). The work shifted from mind-numbing triage to high-value investigative analysis.
- **Regulatory Response:** Regulators were initially cautious about the idea of “AI replacing humans,” but Bank-F presented data from the pilot showing improved results. The outcome – a compliant, faster, and more accurate process – won them over. The bank also demonstrated robust controls like periodic reviews of the model's decisions. Ultimately, Bank-F set an example in the region for modernizing sanctions screening, balancing efficiency with stringent compliance.

Use Case 9: Detecting Trade-Based Money Laundering with OCR Miner

Problem Definition: Bank-G, specializing in trade finance, was concerned about **Trade-Based Money Laundering (TBML)**. Criminals can use falsified trade invoices, over- or under-invoicing goods, and shell import/export companies to move illicit funds under the guise of trade. These schemes are hard to catch with traditional transaction monitoring alone because the red flags lie in the documents (invoices, bills of lading) and not just the payment amounts. Bank-G needed a way to systematically examine trade documents for **sanctions violations, dual-use goods, and value discrepancies** that indicate TBML.

Solution Steps:

1. **OCR Miner Integration:** Bank-G deployed H3M's **OCR Miner** solution into its trade finance processing. When trade documents (e.g., Letters of Credit applications, invoices, shipping documents) are submitted, digital copies are fed into the OCR Miner system.
2. **Optical Character Recognition:** OCR Miner first **ingests and preprocesses** each document (PDFs, scans). It applies optical character recognition to extract text – item descriptions, quantities, values, company names, ports of loading, etc. – converting unstructured paperwork into structured data.
3. **Content Analysis and Vectorization:** The extracted data is then analyzed by the AI:
 - **Keyword/Pattern Matching:** It looks for keywords that indicate risk. For example, names of **dual-use goods** (items with civil and military use) or known **sanctioned commodities** are flagged. In one case, an invoice for “precision machinery parts” raised a flag because the description matched components often controlled for weapons programs.
 - **Value Anomaly Detection:** OCR Miner compares invoice values against market prices. It detected cases of significant over-invoicing (goods priced 5-10x typical market rate) – a classic TBML indicator where extra value transfers hidden funds.
 - **Consistency Checks:** It checks that details make sense – e.g., a bill of lading claiming to ship heavy machinery via an airline (inconsistency), or repeated identical invoice numbers – which could signal fabricated shipments.
 - **Vector Analysis:** Using machine learning, it forms a “vector” profile of each document (a numeric representation of content) and compares it to known legitimate and fraudulent documents. Documents too similar to known suspicious patterns get higher risk scores.
4. **Risk Scoring & Alerting:** For each transaction's document set, OCR Miner generates a **risk score and findings**. For example: *Invoice #123: Score 0.9 (High Risk) – Commodity “Xylene” flagged (dual-use chemical); Invoice value 300% above market; Exporter company on watchlist.* These findings are compiled into an alert sent to Bank-G's compliance team.
5. **Analyst Review:** Analysts receive the OCR Miner alerts alongside the extracted data and document images. In one instance, they saw the tool flagged an exporter who appeared in the Panama Papers leak, and an invoice for agricultural equipment priced exorbitantly. With this info, they investigated further – contacting the client for explanations, cross-checking the companies involved – which ultimately led to blocking the transaction and filing a SAR for TBML suspicion.

Tools Used:

- **OCR Miner:** Advanced OCR and AI module for document analysis in compliance. Bank-G utilized it to detect **sanctions breaches, dual-use goods, and TBML in trade documents**. Features:
 - **Document OCR & Parsing:** Converted trade documents into machine-readable text with high accuracy, enabling further analysis.
 - **Risk Indicators Database:** Built-in knowledge of risky goods (e.g., certain chemicals, electronics), sanctioned entities, typical price ranges, etc., to identify red flags in documents.
 - **Machine Learning Analysis:** Beyond keyword matching, used ML vector analysis to catch patterns of fraud in document sets (like repetitive language or structural anomalies common in fake invoices).
 - **Results & Reporting:** Produced a structured report for each transaction with identified issues and an overall risk score. This report was directly used in Bank-G's case management and SAR drafting, saving analysts time.
- **Trade Data Integration:** The system was linked with Bank-G's trade finance platform, so it automatically fetched relevant reference data (e.g., known market prices from databases, country sanction lists) to enhance its analysis.
- **User Feedback Loop:** The module includes a feedback mechanism where analysts mark alerts as true or false. This trains the AI further (continuous improvement).

Outcome Achieved:

- **TBML Schemes Uncovered:** Bank-G successfully identified multiple TBML schemes. In one case, OCR Miner revealed that a customer was **shipping trivial products (e.g., second-hand clothes) invoiced at millions of dollars**, a strong indicator of money laundering. Another case flagged a **sanctioned country** involved via an intermediary – a container routed through a third country but ultimately destined for North Korea, caught by the system parsing shipping documents. These would have likely gone unnoticed without document analysis.
- **Enhanced Compliance on Trade:** The bank's TBML detection capability went from reactive (relying on external tips or audits) to proactive. They intercepted illicit trades in real-time. Regulators were impressed; during an inspection Bank-G demonstrated how OCR Miner **targets dual-use goods and TBML** risks systematically. This put the bank ahead of peers in trade AML controls.
- **Reduced Manual Review:** Previously, compliance officers had to manually read and interpret lengthy trade documents – a slow and error-prone process. Now, OCR Miner does the heavy lifting, highlighting the suspicious parts of text. This **improved efficiency and accuracy**. The team estimated a 60-70% time savings per document set reviewed.
- **No Missed Sanctions/Export Violations:** With OCR Miner, Bank-G is confident it isn't inadvertently financing prohibited trade. The tool caught sanction-relevant details (like sanctioned entities hidden behind trading companies, or prohibited goods) that transactional screening alone might miss. This drastically lowers the risk of regulatory breaches and fines.
- **Holistic View of Trade Transactions:** By combining transaction data and document insights, the bank gets a 360° view of trade activity. In fact, the insights from OCR Miner feed into their KYC profiles – e.g., a customer engaged in over-invoicing is now flagged high-risk in their system. This cross-pollination of information strengthens overall AML efforts.

Use Case 10: Refining Customer Risk Scoring with KYC Miner

Problem Definition: CU-1 Credit Union relied on a basic KYC risk rating (scoring customers as Low, Medium, High at onboarding) that often became outdated. It didn't account for ongoing behavior or new risk data. As a result, some **high-risk customers were hidden among "Medium"**, and some low-risk customers kept getting flagged by transaction monitoring due to one-off large transactions. CU-1 wanted to implement a dynamic, data-driven **Customer Risk Scoring** system that continuously adjusts with each customer's activities and external risk indicators.

Solution Steps:

1. **KYC Miner Implementation:** CU-1 introduced H3M's **KYC Miner**, an advanced customer risk scoring engine. They fed it initial data on all customers – KYC profiles, account types, historical alerts, transaction patterns – to train an AI model that understands what combinations of factors lead to truly higher risk.
2. **Feature Integration:** The KYC Miner aggregated a wide range of features for each customer:
 - **Static KYC Data:** Country of residence, occupation, product types, presence on PEP/sanctions lists (if any).
 - **Behavioral Data:** Number of alerts triggered in last 12 months, average and max transaction amounts, use of high-risk corridors (e.g., frequent remittances to high-risk countries), sudden changes in activity.
 - **Relational Data:** Connections to other customers (shared addresses, phone, employer) – potentially indicating rings or networks.
 - **External Risk Events:** Any adverse media hits from RUMI or updates like becoming a PEP.
3. **Machine Learning Risk Model:** The KYC Miner used machine learning to weigh these factors. It learned from past cases – for example, customers who ultimately had SARs filed were tagged as high risk in training, whereas long-time customers with no issues served as low-risk examples. The model discerned patterns (e.g., a customer with medium initial KYC risk but multiple international wire alerts is likely higher risk than their initial rating suggests).
4. **Continuous Scoring & Alerts:** CU-1 set the system to **recalculate risk scores monthly** for each customer, or immediately upon certain triggers (like a new SAR filed or a major adverse media event). The output was a updated **risk rating (score 0-100 and Low/Med/High category)** for every client. Notably, about 5% of their customers moved to a higher risk tier thanks to behavioral factors. Whenever a customer's risk tier jumped significantly, an alert was sent to compliance to review that account.
5. **Dashboard & Drill-Down:** The results were made visible via a **KYC 360° Dashboard**. For a given customer, analysts could see an integrated profile: initial KYC info, current risk score, and contributions of various factors (e.g., "High risk due to: 3 overseas transactions to high-risk country X + shared address with known fraud suspect"). They could drill down on alert history and related entities in one view.

Tools Used:

- **KYC Miner:** AI-driven customer risk scoring and management tool. It provided CU-1 with:
 - **Holistic Risk Assessment:** Combined static KYC data with dynamic inputs (transactions, alerts, external data) to produce an all-encompassing risk profile per customer.
 - **Automated Risk Updates:** The system automatically re-scores customers on schedule and upon key events, ensuring risk ratings are **never stale**.
 - **Integration with Alerts & External Data:** Directly pulls in data from transaction monitoring, sanctions screening, and adverse media (RUMI), aligning with a 360° risk approach.

The screenshot shows the 'KYC Scenario Miner' interface. At the top, there's a 'SELECT SCENARIO' section with a search bar and a list of scenarios: 'KYC_Webinar_1' (selected), 'glam_kyc_msc Register', 'TestKycMiner1', and 'test-said'. To the right, there are input fields for 'Very High Risk Threshold' (100), 'High Risk Threshold' (80), and 'Medium Risk Threshold' (60). Below these, there's a 'DEFINITION' field with a placeholder text 'KYC Miner use for the webinar'. On the right side, there's a 'STATUS' dropdown set to 'Passed', and fields for 'CREATED DATE' (2025-01-20), 'MODIFIED DATE' (2025-01-20), and 'MODIFIED BY' (edga). Below this, there's a table with columns: 'Feature Name', 'Operator', 'Feature Details', 'Score', and 'Action'. The table contains five rows of data for features like BIRTHDATE, CREATEDAT, INCOME, NATIONALITY, and POPULATIONGROUP. At the bottom, there are 'SAVE' and 'SAVE AS' buttons.

Outcome Achieved:

- **Dynamic Risk Visibility:** CU-1 now has a **living risk profile** for each member. Instead of one-and-done risk classification at onboarding, they can see risk evolve. This meant that when a normally medium-risk customer started frequent large cash deposits and international transfers, the system upgraded their risk in real-time and compliance intercepted the issue.
- **Earlier Detection of High-Risk Customers:** Several customers who would have remained under the radar as “Medium Risk” were elevated to “High” by KYC Miner due to their behavior. For instance, one small business owner was re-scored high after funneling funds to a high-risk country; EDD revealed the business was likely a shell. These proactive flags helped CU-1 take action (SAR or exit) **months earlier** than under the old system.
- **Fewer Missed Red Flags:** Conversely, the improved profiles reduced cases of ignored risk. If a customer had multiple smaller red flags across different systems (a mild sanction hit here, a few alerts there), KYC Miner aggregated those to show the customer as high risk overall. This **integrated view of CDD, sanctions, and transaction monitoring data provided all-encompassing risk profiles** and ensured no combination of risk factors was overlooked.

Use Case 11: Holistic Risk Oversight with KYC 360° Dashboards

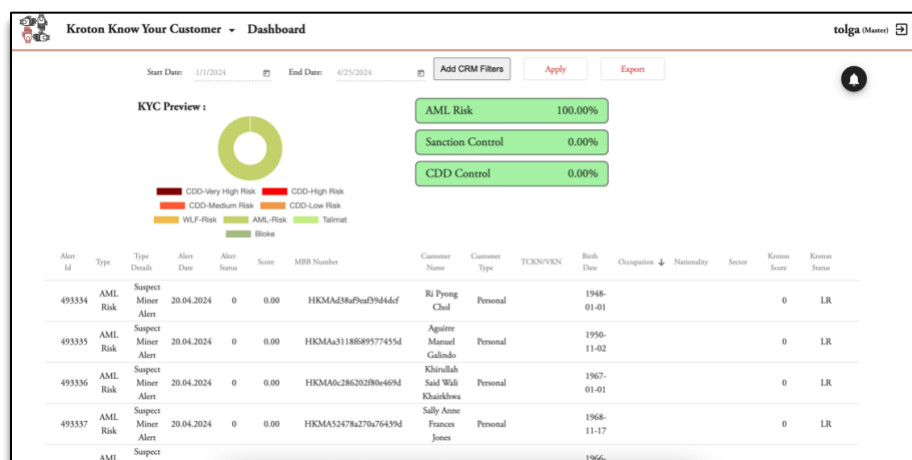
Problem Definition: Bank-H's compliance data was fragmented. KYC reviews, transaction alerts, sanction hits, and case notes were stored in different systems. This made it difficult for compliance officers to see the **full picture of a customer's risk** or to answer questions like “who are our riskiest customers?” effectively. They wanted a **unified dashboard** that would bring together all compliance data (CDD, transactions, sanctions, adverse media) into a comprehensive, easily digestible view for both analysts and management.

Solution Steps:

1. **Data Integration:** Bank-H deployed H3M's **KYC 360° Dashboards** solution. The implementation team connected various data sources to the dashboard: the KYC database (with customer profiles and documents), the transaction monitoring system (alerts and outcomes), the sanctions screening log (any hits or cleared matches), and even the RUMI adverse media search results.
2. **Dashboard Configuration:** They designed the dashboard with multiple views:
 - **Customer Risk Overview:** A table listing all customers with their current risk scores or tiers, sortable by risk. High-risk customers are highlighted.
 - **Customer Drill-down Page:** Clicking a customer brings up a detailed profile: KYC details (onboard date, KYC risk rating, PEP status), a summary of recent alerts (e.g., 3 alerts in last 6 months, with statuses), any sanctions or PEP hits, linked accounts, and recent adverse media hits if any. It also shows the trend of their risk over time and comparisons to peer group.
 - **Aggregate Risk Metrics:** Charts showing distribution of risk levels across the bank, top 10 riskiest customers, risk by geography (map of customer locations color-coded by risk), and by business line (e.g., corporate vs retail).
3. **User Access:** Compliance analysts and investigators use the dashboard daily. For example, when assigned an alert on a customer, an investigator opens that customer's 360° profile. They can quickly see if this customer has a history of alerts or other red flags. Recently, an analyst reviewing an alert saw on the dashboard that the customer had a prior SAR and was related to another account that was under investigation – context that was crucial and immediately available via the dashboard.
4. **Group and Geo Analysis:** The compliance managers utilize the group insights. They filter the dashboard to see risk by region. The dashboard revealed, for instance, that **customers from X branch had disproportionately high alerts**, prompting a targeted review of that branch's onboarding processes. They also used it to identify if certain **customer segments** (like money service businesses) were driving risk, which guided policy adjustments.
5. **Reporting & Communication:** Bank-H leveraged the dashboards in reporting to senior management and regulators. In quarterly risk committee meetings, the MLRO shows snapshots from the dashboard: e.g., “We have 120 High risk customers, here's the breakdown by type and the actions taken on them.” If needed, they can drill in live to answer board member questions, a big improvement over static spreadsheet.

Tools Used:

- **H3M KYC 360° Dashboard:** A comprehensive compliance risk visualization and management tool. It provided:
 - **Unified Risk Profiles:** Integrated data from multiple dimensions (CDD, transactions, sanctions, media) to present an **all-encompassing view of each customer's risk**. This eliminated the need to log into five different systems to gather info.
 - **Interactive Drill-Down:** Users can click on a customer to see full historical details and risk factors, or click on a chart segment to filter (e.g., see all customers in a high-risk country). It's a **one-stop investigative interface** for analysts, with search and filter functions to quickly find any customer or cohort.
 - **Real-Time Updates:** The dashboard data updates as underlying systems update. If a new alert is filed, it reflects in the customer's profile immediately. If a customer's status changes (e.g., added to a watchlist), the dashboard highlights it. This real-time aspect means decisions are based on current info, not last month's data.



Outcome Achieved:

- **Holistic Insight:** Bank-H finally has a **360° view** of customer risk at their fingertips. This led to better decision-making. For instance, in one case a customer's transaction alert seemed minor, but the dashboard revealed the customer also had negative news and a past fraud alert – together painting a picture of significant risk, prompting a SAR.
- **Efficiency and Speed:** Investigations sped up by an estimated 30%. Analysts no longer scramble for information – everything is on one screen. This saved precious time, especially in complex cases. One investigator said it's *"like having an instant briefing on the customer"* before diving into the specifics of an alert.
- **Improved Communication:** Regulators who examined the system were satisfied that Bank-H can *"examine the complete risk landscape with integrated views of CDD, sanctions, and transaction monitoring"*. This demonstration of control and oversight led to a smoother regulatory exam with few findings.

- **Elevated Compliance Program:** The **comprehensive risk visualization and management** means they allocate resources smarter, identify issues faster, and can demonstrate control to stakeholders.

Use Case 12: Uncovering Hidden Risks with RUMI

Adverse Media Search

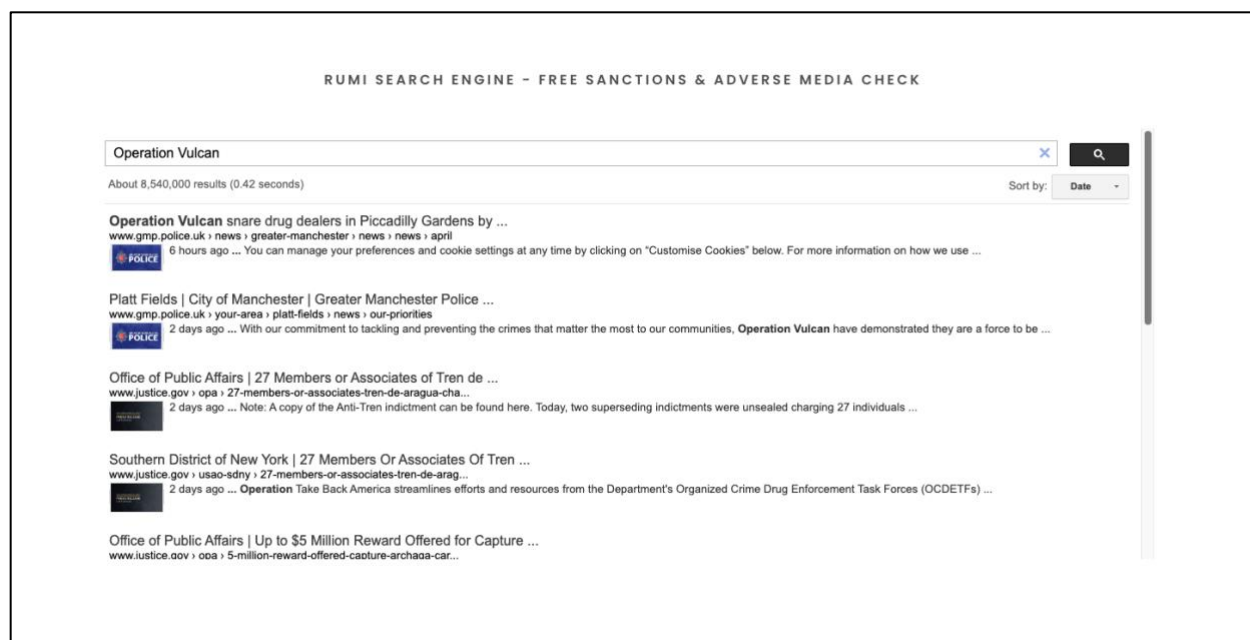
Problem Definition: MSB-3, a digital payments provider, must vet customers and counterparties for any adverse media or negative information that might indicate risk (e.g., involvement in fraud, crime, or being on watchlists not easily visible). Relying on Google searches was inconsistent and often missed deep-web information like **offshore leaks or foreign-language news**. MSB-3 feared that without a specialized search tool, they could be onboarding clients with serious hidden red flags or missing crucial context during investigations.

Solution Steps:

1. **Adverse Media Tool Adoption:** MSB-3 started using **RUMI** – H3M’s compliance-focused search engine – as part of both onboarding due diligence and investigation processes. Analysts were trained to run every new high-risk customer name and any suspicious entity through RUMI.
2. **Comprehensive Search Query:** When an analyst inputs a subject (individual or company) into RUMI, the tool’s **Compliance-Autocomplete** suggests relevant expansions. For example, typing a company name might prompt “<Name> + money laundering” or “<Name> + fined” to ensure thorough search terms. Analysts can select suggestions or enter their own.
3. **Specialized Sources Scanned:** RUMI then searches across a broad array of sources tailored for compliance:
 - o Global and local news outlets, including ones in high-risk regions (with translation as needed).
 - o Regulatory and legal databases (e.g., enforcement actions, court cases).
 - o Leaked data sets like the **Panama Papers, Paradise Papers**
 - o Cryptocurrency wallet blacklists (if applicable) and deep web forums.
 - o Sanctions and watchlists (as an extra layer to their sanctions screening).
 - o Proprietary H3M databases of PEPs and adverse profiles.
4. **Prioritized Results:** In seconds, RUMI returns results ranked by relevance using advanced algorithms. It **prioritizes compliance-relevant hits** – e.g., a small local news article about a fraud involving the person will rank higher than an unrelated social media mention. Each result shows a snippet of context around the search term.
 - o In one case, searching a new corporate client’s name immediately surfaced a foreign news piece about that company being fined for illegal gambling operations – something not found on Google.
 - o Another search connected an individual to a shell company in the **Panama Papers database**, revealed through RUMI’s integrated leak search.
5. **Deep Dive and Monitoring:** Analysts click on relevant results to read full details. RUMI’s interface allows filtering by source type (e.g., show me just news vs. just leaks). MSB-3’s team set up alerts in RUMI for certain high-profile names so that if new info appears (like a news article tomorrow), they get notified. This was especially used for their PEP customers.
6. **Network Analysis:** As a follow-up, RUMI discoveries were fed into Link Miner for a graph analysis, revealing a money laundering network, that otherwise would not be visible to any scenarios or machine-learning models.

Tools Used:

- **RUMI Search Engine:** A compliance-tailored adverse media and data search tool. MSB-3 leveraged its unique capabilities:
 - **Compliance-Autocomplete**
 - **Extensive Data Sources**
 - **Relevance Ranking:**
 - **Advanced Filters**
- **Integration with Link Miner:** Link Miner discovered hidden links with other MSB-3 clients by using advanced network analysis techniques.



Outcome Achieved:

- **Deeper Due Diligence:** MSB-3 significantly improved the depth of its KYC investigations. In at least a dozen cases in the first quarter, RUMI uncovered **critical adverse information** that was missed by earlier processes. These included ties to criminal entities, involvement in lawsuits/fraud, or appearances in leaked documents.
- **Faster Investigations:** What used to take hours of manual searching across Google, sanctions lists, and various registries is now accomplished in minutes with a single RUMI query. Analysts report that for complex cases, they save several hours and get more reliable results.
- **Improved Risk Decisions:** In one scenario, multiple customers were sending funds to the same obscure business. RUMI revealed that business had been mentioned in a forum as a front for illegal online pharmacies. The phrase *“several flagged entities were linked to shell companies and had adverse media associations”* became a reality that RUMI highlighted clearly.
- **Regulatory Kudos:** Regulators often emphasize adverse media checks. MSB-3 was able to demonstrate a robust process using RUMI, showing logs of searches and findings for higher-risk clients. This proactive approach exceeded expectations. Regulators were particularly pleased that MSB-3 could find hard-to-get information and discovered the full extend of a very important network.

Use Case 13: Community Approach – Fighting Financial Crime via Federated Learning

Problem Definition: A group of five small financial institutions suspected that criminals were exploiting the lack of visibility across institutions – e.g., orchestrating fraud or money laundering by spreading transactions and mule accounts among them so no single bank’s monitoring would catch the pattern. However, directly sharing customer data between banks wasn’t feasible due to privacy and regulatory constraints. They needed a way to **learn from each other’s AML data without exposing that data**, to collectively detect cross-bank money laundering schemes.

Solution Steps:

1. **Consortium Formation:** The five banks formed a consortium and enabled H3M’s **Federated Learning** capability within their KROTON platforms. Each bank had its own Suspect Miner model. Under federated learning, these models can **collaboratively train** with others without exchanging raw data.
2. **Local Model Training:** Initially, each bank trained a Suspect Miner on its own transaction data and alert history. For example, Bank A’s model learned patterns in Bank A’s data (like common behaviors of its bad actors), Bank B’s model on Bank B’s data, and so on.
3. **Federated Model Exchange:** At agreed intervals (monthly), the banks’ systems sent the **encrypted model parameters** (not actual transactions or PII) to a central server. H3M’s federated server aggregated these parameters – essentially averaging the learned patterns from all banks, weighted by data volume. The aggregated model represented a more comprehensive knowledge of suspicious patterns, as it had effectively “learned” from combined experience (e.g., a pattern seen at Bank B and C).
4. **Global Model Distribution:** The server then sent back the updated global model parameters to each bank’s Suspect Miner. Each bank updated its model with these new parameters, **immediately enhancing its ability to detect patterns originally seen at other banks**. For instance, Bank A’s model might now catch a behavior that only Bank C had seen before (and vice versa).
5. **Pattern Detection:** Soon after, cross-institution patterns began to surface in alerts. One success: a money mule network where mules at different banks were depositing funds and then funneling to each other in a loop. Individually, each bank saw only pieces (one mule with some unusual deposits). After the federated update, their models recognized the larger pattern – Bank A flagged its mule as suspicious because the profile now matched what was seen at Bank B and D. Each bank’s alerts pointed to counterparts in other banks.
6. **Collaboration on Investigations:** With these alerts, the banks legally shared just the necessary info via their consortium agreement (e.g., Bank A informs Bank B of an account number that its model flagged as likely part of a cross-bank scheme). They confirmed the interconnected transactions forming the network. Law enforcement was brought in with a combined SAR that mapped out the full scheme across banks.
7. **Iterative Improvement:** Every month, this federated learning cycle repeats. Each cycle further refines the joint model as more feedback (like confirmed SARs) is included. The consortium expanded to include a couple more regional banks, further strengthening the model with more data.

Tools Used:

- **H3M Federated Learning Framework:** Enabled **multi-institution model training** without sharing sensitive customer data. Key aspects:
 - **Secure Parameter Sharing:** Only model weight updates (essentially abstract numbers) are shared and aggregated. No raw transactions or customer identities leave a bank, preserving privacy and compliance with data protection laws.
 - **Central Aggregator:** A server that takes in all participating models' parameters and produces a unified model. This ensures **collaborative improvement of detection models** across banks.
 - **Federated Updates:** Regular distribution of the improved model back to each bank's Suspect Miner. The process is seamless and encrypted end-to-end.
- **Suspect Miner (at each bank):** The local ML model that now benefits from a **much larger effective training dataset** (combined knowledge of all banks). Each still runs independently on local data to generate alerts, but its brainpower is multiplied by federation.
- **Consortium Governance:** While not a software tool, it's worth noting the banks established legal agreements and a committee to oversee the federated learning collaboration – ensuring compliance and deciding on model update frequency, etc. H3M provided reporting metrics to this committee, like improvement in detection rates attributable to the federated model.

Outcome Achieved:

- **Cross-Bank Schemes Detected:** The community banks collectively identified at least three major multi-bank money laundering networks in the first year. One example: a syndicate was cycling illicit funds through accounts at Banks A, C, and E in a round-robin fashion to evade singular detection. The federated-enhanced models at each bank flagged the activity, and when the banks compared notes through the consortium, the full picture emerged. Such a scheme **would likely have gone unnoticed** without this collaborative approach.
- **Improved Individual Bank Detection:** Each bank saw a notable uptick in detection performance. True positives increased because models became smarter with more data; false positives in some cases decreased, as the model learned more differentiating factors from the collective data. For instance, Bank D started catching suspicious activity involving cryptocurrency that it had no experience with prior, because the model learned from Bank B (which had faced crypto scams).
- **Privacy Preserved:** Importantly, this was achieved **without sharing customer PII or transactions between banks**, satisfying regulatory and privacy concerns. Regulators were intrigued and supportive when shown how federated learning allowed this outcome – essentially enabling an “industry-level brain” to fight crime without a central data pool. This pioneering approach was highlighted as a future model for AML collaboration.
- **Scalable Network Effects:** As more institutions join the federated network, everyone's models get even better. The community banks demonstrated how even smaller banks can achieve **big-data AI advantages** by banding together. Over time, the network effect could make it very difficult for criminals to exploit gaps between banks, as a pattern learned in one place immediately becomes known in all.

Use Case 14: Multi-Module AI Workflow Stops Complex Money Laundering Scheme

Problem Definition: MSB-4, a global e-money and remittance company, faced a complex money laundering challenge: illicit funds were moving through customer accounts in patterns that spanned multiple channels and were hard to detect with any single measure. The scheme involved **high-risk customers** hidden among millions of users, **interlinked transactions** that only made sense when viewed as a network, and use of front companies. Traditional methods produced too many false positives or missed the subtle coordination. MSB-4 needed an integrated approach leveraging all of H3M's modules to detect and evidence this scheme end-to-end.

Solution Steps:

Step 1: Pinpointing Suspicious Entities – Suspect Miner

MSB-4 deployed **KROTON Suspect Miner** on its transaction dataset to identify suspicious customers that scenario-based rules hadn't flagged. The AI model, using over 100 features (transaction patterns, location, device info, etc.), produced a **curated list of high-risk entities with over 85% true positive accuracy**. For example, it highlighted a cluster of customers who, while not triggering existing rules, had unusual patterns (many small incoming transfers followed by large outbound remittances). These became the starting suspects for the investigation.

Step 2: Unraveling the Network – Link Miner

Next, MSB-4 used **Link Miner** to map the relationships among the flagged suspects. By inputting those high-risk customers, Link Miner constructed a network graph of transactions and shared attributes. This revealed an **intricate money flow network** – certain suspects were intermediaries linking many senders to a few receivers. The visual graph instantly made clear a structure that would have taken months to piece together manually (one node was central, receiving funds from multiple flagged senders). Link Miner's fast algorithm handled the millions of nodes MSB-4 has, producing results in under an hour. The output identified key "hub" accounts that appeared to be controllers of the laundering ring.

Step 3: Comprehensive Background Checks – RUMI Search Engine

With a network in hand, MSB-4 investigators performed deep background checks on the entities involved using **RUMI**. They searched the names of the central individuals and entities uncovered by Link Miner. RUMI pulled from global databases: it found, for instance, that one shell company in the network was mentioned in the Panama Papers leak and had adverse media linking it to a corruption scandal. Another individual was found on a watchlist forum as a known money mule recruiter. These findings provided **contextual evidence**: several flagged entities were linked to shell companies and adverse media, corroborating that the network was engaged in illicit activity.

Step 4: Holistic Risk Synthesis – KYC 360° Dashboard

MSB-4's team compiled all insights in a **360° dashboard** for decision-makers. Each suspect's profile was updated with their Suspect Miner risk score, Link Miner connections (displayed as a sub-graph around them), any sanction/PEP hits, and the adverse media findings. This holistic view made it easy for compliance leadership to grasp the full picture: who the players were, how they were connected, and what external evidence supported the suspicion. The dashboard showed, for example, that the top 10 suspects were interrelated and collectively had moved \$5M through the network.

Tools Used:

- **Suspect Miner:** AI module that identified unknown high-risk customers with minimal false positives, kicking off the investigation with a focused list.
- **Link Miner:** Network analysis tool that exposed the **full structure of the laundering network**, turning disparate alerts into a cohesive story of illicit fund flows.
- **RUMI Adverse Media Search:** Provided external intelligence (shell company ties, media reports) that strengthened the case by showing links to known illicit activities.
- **KYC 360° Dashboard:** Integrated all data from the above modules into a single investigative and reporting interface, enhancing clarity for decision-makers.

Outcome Achieved:

- **No Stone Unturned Detection:** The multi-module approach ensured that **every facet of the scheme was uncovered** – from identifying key suspects with AI to mapping their network and finding corroborating evidence. MSB-4 not only caught the launderers but understood their methods, allowing for prevention of similar schemes.
- **High Efficiency and Accuracy:** By leveraging each tool for what it does best, MSB-4 achieved in days what could have taken months. The process was largely automated and data-driven, resulting in minimal false leads. Each alert that passed through the funnel was highly credible. This efficiency is evidenced by the fact that each module's output was crucial and **pivotal, streamlining what would traditionally be a long-winded process.**
- **Successful Disruption of ML Network:** The identified network was effectively dismantled – accounts were frozen and the perpetrators were reported to law enforcement. Subsequent monitoring showed that those patterns ceased within MSB-4's ecosystem. It's likely that without this integrated approach, the network would have continued operating undetected.
- **Persuasive Reporting:** The compiled case with network graphs and media evidence made for a compelling SAR. It demonstrated MSB-4's thoroughness and helped authorities quickly grasp the situation, increasing the chance of enforcement action. Regulators, in turn, praised MSB-4 for the depth of analysis. The quote “**no stone was left unturned**” could well describe the regulators' view of MSB-4's investigation.

Use Case 15: Multi-Module Crackdown on Trade-Based Money Laundering

Problem Definition

Bank-J's trade-finance desk flagged a growing set of oddities: customers over- or under-invoicing goods, repetitive round-dollar payments, and shipping documents that looked "too generic." These clues sat in different systems—transaction monitoring, trade-finance operations, and KYC files—so no single view proved a scheme. Compliance leaders suspected a coordinated TBML network and decided to weave insights from every advanced module into one investigation.

Solution Steps

1. **Suspect Miner – Transaction Anomaly Detection**
 - The model sifted three years of trade-related SWIFT and domestic wires.
 - It highlighted an electronics exporter receiving identical \$500 k payments every 14 days from shell importers—classic over-invoicing.
 - Output: 37 customers and 312 payments tagged *high-risk TBML*.
2. **Link Miner – Mapping Shell Networks**
 - Analysts seeded Link Miner with the 37 suspects.
 - Graph analysis showed five exporters and three importers shared directors, phone numbers, or circular fund flows that converged on one offshore account in Curaçao.
 - Result: a single, eight-node laundering ring across three jurisdictions.
3. **OCR Miner – Trade Document Forensics**
 - Pulled invoices, packing lists, and bills of lading for the 312 payments.
 - OCR flagged:
 - "Industrial machines" priced 4-6× market rate.
 - Reused invoice numbers and mismatched container weights.
 - Dual-use chemicals shipped to a high-risk port without export license codes.
4. **Sanctions Miner – Export-Control & Vessel Screening**
 - None of the parties were sanctioned, but Sanctions Miner scored the commodity (precursor chemicals) and the vessel *Sea Glory*—previously cited for Iran shipments—at elevated risk.
 - Added a sanctions-compliance dimension and urgency to the case.
5. **RUMI – External Adverse Media & Leaks**
 - Searches linked one exporter to a VAT-fraud probe in Poland and named an owner in Paradise Papers data.
 - Route-level search uncovered prior cocaine seizures on the exact shipping corridor.
6. **Case Assembly – KYC 360 Dashboards & SARs**
 - Dashboard stitched all inputs into a timeline: anomalous payments → shared ownership → forged docs → high-risk goods → adverse media.
 - Bank froze transactions, filed comprehensive SARs with network graphs and sample invoices, and alerted customs and FIUs in three countries.

Tools Employed

- **Suspect Miner** – pinpointed repetitive high-value payments.
- **Link Miner** – visualized cross-border ownership & flow links.
- **OCR Miner** – exposed forged or inflated trade documents.
- **Sanctions Miner** – flagged controlled goods and watch-listed vessel.
- **RUMI** – supplied press and leak evidence connecting actors to prior fraud.
- **KYC 360 Dashboard** – unified intel for rapid executive sign-off.

Outcomes

- **Network Disrupted:** \$40 million in planned TBML flows halted; detailed evidence pack sent to law enforcement.
- **Regulatory Commendation:** Supervisors praised Bank-J’s “integrated AI/OCR controls” for trade finance—rare in the market.
- **Enhanced Controls:** Bank now mandates OCR Miner and Link Miner review for any high-value trade payment flagged by transaction monitoring.
- **Efficiency:** Integrated workflow let eight investigators handle a complex, multi-jurisdiction case without massive staffing, setting a repeatable playbook for future TBML threats.