HAWK

Whitepaper

Precision at Scale: Reimagining Sanctions & Watchlist Screening with Hybrid AI Frameworks

How Pairing Contextual Machine Learning with AI Agents Reduces Alert Noise, Scales pKYC, and Strengthens Screening Governance



HAWK

Contents

Executive Summary 3

The State of Screening: Dual Pressures on Payments & Customer Profiles 4

The AI Adoption Gap in Screening Today 5

The Business Impact & ROI for AI in Screening 6

Hybrid AI Framework for Screening 7

How to Evaluate Screening AI Solutions 12

Decision Governance: Straight-Through-Processing 15

Keeping Costs Predictable at Scale 16

Hawk's Approach: Hybrid AI That Improves Over Time 17

Future Outlook 21

Conclusion & Next Steps 22

© 2026 Hawk or its affiliated companies. All rights reserved

Executive Summary

Financial institutions face unprecedented operational pressure in sanctions and watchlist screening. Expanding global sanctions lists, strict liability enforcement, and the industry transition toward continuous Perpetual KYC (pKYC) re-screening have flooded compliance teams with false positive alerts. Fuzzy matching engines–the backbone of screening–generate 90–95% noise, draining analyst capacity and causing severe payment delays.

Yet despite the scale of this challenge, AI adoption in screening remains surprisingly limited: research from Chartis and Hawk found that only 16% of institutions use AI at scale across their screening programs, leaving most compliance teams to resolve false positives through manual review alone. Closing this gap requires a hybrid AI framework that operates at both ends of the screening lifecycle–contextual machine learning that derives missing data points and suppresses false hits before they are ever created, paired with autonomous AI agents that investigate, triage, and close the alerts that remain. Deployed together, these two layers turn screening from a function that scales with headcount into one that scales with technology.

Where to Start

- **Resolve missing data before it creates alerts.** Contextual machine learning works out details like gender and entity type in real time. That removes up to 25% of alert noise.
- **Bring in AI agents for the alerts that remain.** They gather the evidence, check it against your policy and document the rationale in under a minute.
- **Keep payments moving while agents go deeper.** Machine learning screens in milliseconds before the block, so settlement SLAs hold. Agents work after the hit, so the extra depth doesn't slow payments down.
- **Filter before you spend on LLM tokens.** Let machine learning remove false hits before any large language model is called. This keeps run-rate costs to $1 – $5, well below agent-only approaches.
- **Decide your own level of oversight.** Configure straight-through-processing and set confidence thresholds that are allowed by your regulators and that match your risk appetite, with a full audit trail behind every decision.



© 2026 Hawk or its affiliated companies. All rights reserved

The State of Screening: Dual Pressures on Payments & Customer Profiles

Sanctions and watchlist screening operates under strict liability: a single missed true positive match can lead to regulatory fines, severe reputational damage, or loss of banking licenses. In response, compliance teams run broad fuzzy matching nets, which flood operations with coincidental name matches.

Today, compliance teams face twin operational strains: real-time cross-border payment settlement demands sub-second processing, while continuous Perpetual KYC (pKYC) re-screening generates endless background alerts for customers.

Table 1: Payment Screening vs. Customer Screening Approach & Challenges

Dimension	Payment Screening (Sanctions)	Customer Screening (PEPs, Sanctions & Adverse Media)
Trigger Mechanism	Real-time, in-flight payment messages (SWIFT, ISO 20022, FedNow)	Continuous portfolio re-screening & pKYC delta triggers
Operational Impact	Blocked payments, delayed settlements, and customer friction	Massive backlog of non-material background alerts, delayed customer onboarding
Primary Root Cause	Lack of data richness for payment data in scope of screening (no date of birth, geography and other relevant secondary information)	Extensive alerts across different risk types that may lack sufficient secondary data (e.g. adverse media)
Analyst Burden	Hours spent clearing obvious false positives	Hours spent gathering additional information on customer/risk entity

Diagram 1: How fuzzy matching results in 90–95% of hits being false positives

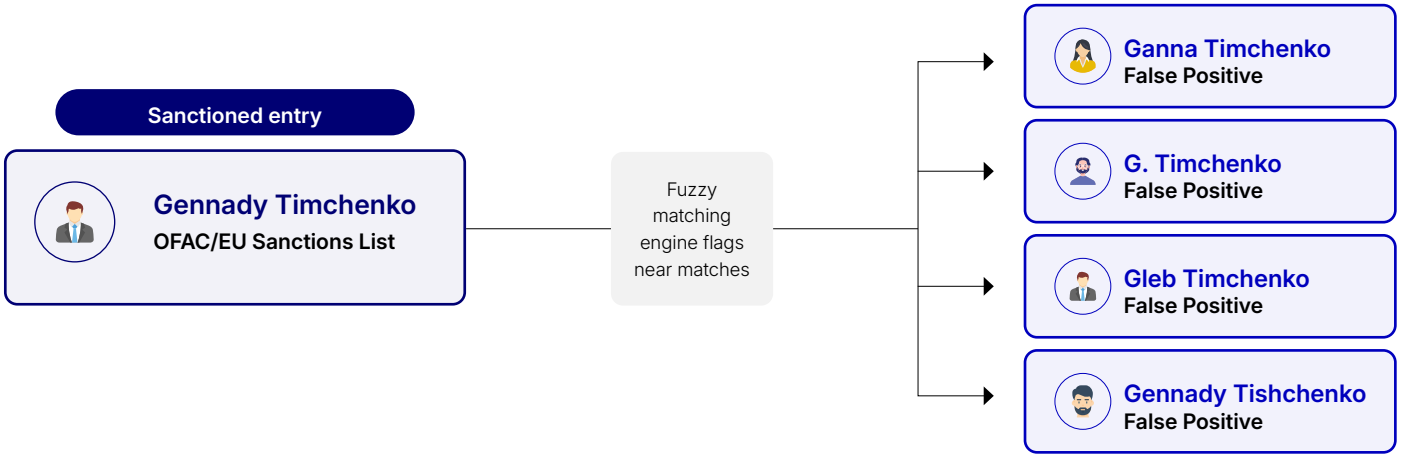

© 2026 Hawk or its affiliated companies. All rights reserved

The AI Adoption Gap in Screening Today

Across the anti-financial crime industry, AI has already won the argument in principle, but it's the execution within screening that is lagging. Independent research from Chartis found that virtually every financial institution now uses AI somewhere in their fraud program, with 100% reporting some level of adoption. Sanctions screening tells a very different story: only 6% of institutions have operationalized AI at scale, and a combined 64% remain stuck testing, piloting, or exploring it rather than running it in production.

That stall has a real operational cost. Separate research from Celent, surveying compliance, risk and operations leaders at US mid-market banks and credit unions, found that ***77% cite analyst staffing shortages as a top-three challenge in their compliance programs,*** and high false positive alert rates follow close behind as a leading operational friction, cited by 40%. Put simply: the volume of noise screening generates is outpacing the number of people available to review it.

The gap isn't for lack of trying. Chartis's research points to two concrete barriers holding institutions back: ***60% rank insufficient internal AI expertise as their top business challenge,*** and 62% cite poor-quality or incomplete training data as their primary technical obstacle. Even compliance teams motivated to modernize are often blocked by these foundational issues.

Despite these barriers, the direction of travel is clear. ***84% of banks plan to increase their investment in agentic AI, and 92% plan to increase generative AI spending,*** over the next two to three years. The institutions surveyed aren't questioning whether to close this gap, only how quickly they can.

Diagram 2: AI Adoption by Compliance Domain (Chartis Research Data, 2025)

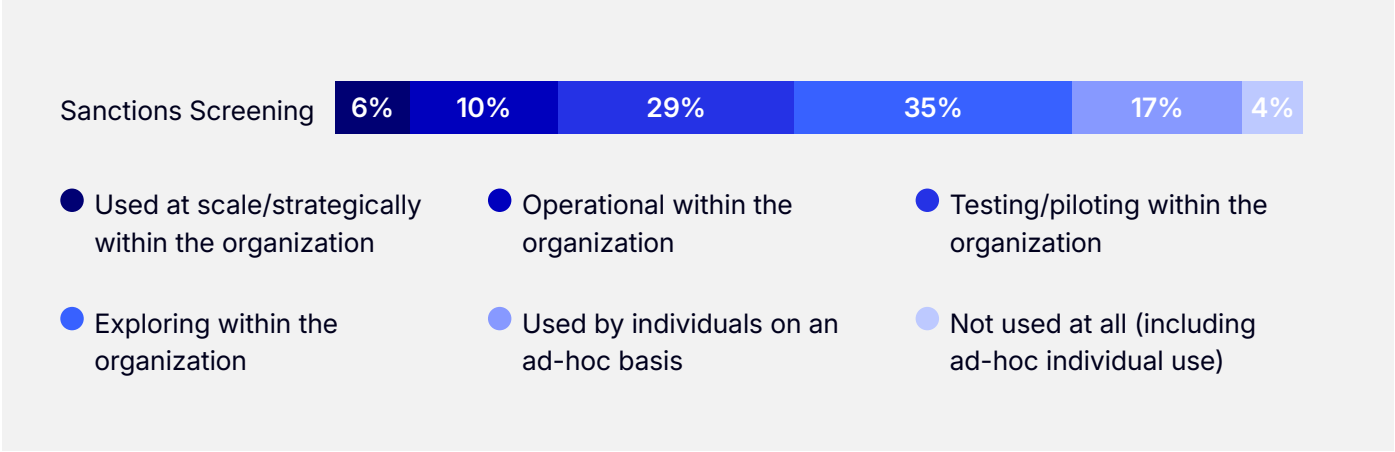

© 2026 Hawk or its affiliated companies. All rights reserved

The Business Impact & ROI for AI in Screening

Applied well, AI in screening can ease some of the operational pressure on compliance teams. The impact shows up in four areas:



Reducing Friction on Payment Settlement
Clearing clear-cut false positives before a payment is blocked preserves the sub-second settlement speed real-time rails demand. Machine learning running in milliseconds lets genuine transactions pass without the customer friction or partner SLA breaches that come from holding payments for manual review.



Increasing Perpetual KYC Velocity
Continuous pKYC re-screening generates a constant stream of background alerts, most reflecting no material change in risk. AI can suppress that noise by cutting clear-cut false positives to let analysts focus on profiles that have genuinely changed. The next page of this whitepaper looks at how.



Standardizing Audit Readiness
Manual review can result in inconsistencies, as analysts under time pressure write up comparable alerts differently. A machine-generated evidence trail replaces that variability with a consistent, defensible record for every decision.



More Predictably Manage Operational Costs
This is where the return is measurable. Chartis research found 48% of institutions using AI in financial crime compliance had already saved more than $1 million in the past year; in AML specifically, 71% had realized cost savings and 94% expect more within 12 months, over half of those expecting more than $5 million annually. Paired with an up-to-80% reduction in manual review workload and an agentic run-rate cost of $1–$5 per alert, the case is no longer theoretical.

Calculating the ROI for Your Institution

For most institutions, the clearest return comes from analyst time. Three inputs shape it: how many false positive hits your team reviews each year, how long each one takes, and what an analyst costs. Together, they show what false positives cost you today. AI improves that figure in two ways: fewer false positives reach an analyst, and the ones that do take less time to resolve. Set against the cost of running the technology and modelled over several years, that gives you a clear picture of your return.

Hawk can model this for you. Share a few details about your screening volumes and team, and we'll build an ROI estimate based on your own operation.

Get in touch here

© 2026 Hawk or its affiliated companies. All rights reserved

Hybrid AI Framework for Screening

As adoption of AI for screening accelerates, the framework that financial institutions choose will matter as much as the decision to invest. To achieve speed, accuracy, and cost governance, modern AI screening relies on a dual-layer hybrid framework.

Two Jobs, Two Kinds of AI

Traditional machine learning is built for speed and scale, running against every transaction or profile in true real time (milliseconds), at a low cost. Agentic AI, orchestrating large language models, is built for depth – it can investigate and gather evidence as an analyst would, but not fast or cheaply.

Diagram 3: How Machine Learning False Positive Reduction and AI Agents Work

Machine Learning
(False Positive Reduction Layer)

- **Fills in missing context:** Predicts gender, entity type, or country of registration directly from a name
- **Compares against the watchlist record:** Checks the newly derived context for contradictions
- **Scores the match:** Produces a 1–100 likelihood score, broken down by attribute
- **Clears or flags:** Closes clear-cut false positives where the institution has enabled straight-through processing and passes ambiguous hits on for review.

AI Agents
(Review Layer)

- **Gathers evidence:** Pulls from internal records and external sources
- **Assesses against policy:** Checks the hit's materiality against the institution's own rules
- **Documents everything:** Logs reasoning, confidence score, and citations into a full audit trail
- **Recommends, escalates or closes:** Recommends a disposition for analyst sign-off, escalates when the evidence isn't conclusive or clears alerts that meet an institution's confidence threshold for straight-through processing

© 2026 Hawk or its affiliated companies. All rights reserved

Layer 1: Closing the Data Gap Before a Hit Opens with Machine Learning AI

SWIFT and ISO 20022 payment messages frequently omit fields like gender, entity type, or jurisdiction – the exact gaps fuzzy matching can't resolve on its own. AI can close them by deriving the missing attribute directly from the name itself:

Image 1: *Example of Machine Learning AI Alert Within the Hawk Platform*

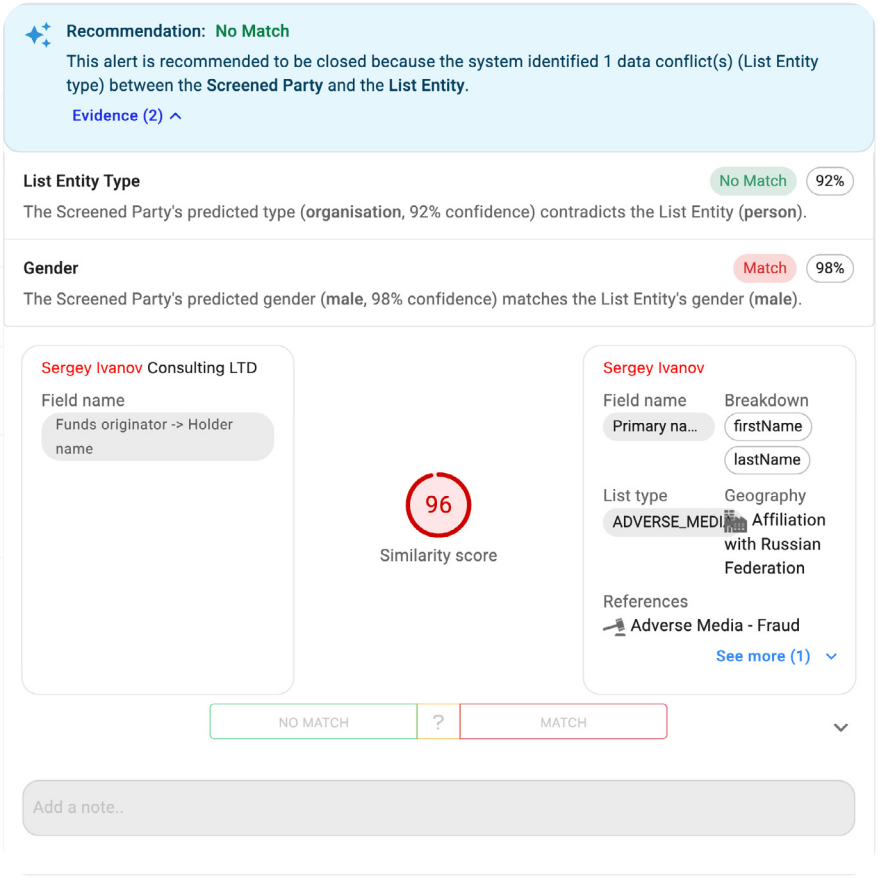

Diagram 4: Machine Learning Contextual Derivation of Missing Data

Gender

Predicts the screened party's gender from their name and compares it to the watchlist entity's.

- **Match:** Michel Jones (male, 99%) vs. watchlist entity Michal Jones (male)
- **No match:** Michelle Jones (female, 97%) vs. watchlist entity Michal Jones (male)

Entity Type

Multilingual language models determine whether a name refers to a person or an organization.

- **Match:** Gazprom Ltd (org, 99%) vs. watchlist entity Gazprom (org)
- **No match:** Nicholas Gazprom (person, 97%) vs. watchlist entity Gazprom (org)

Country of Registration

Detects legal form designations in entity names to infer jurisdiction and rule out country mismatches.

- **Match:** Contoso (German registration) vs. Contoso GmbH
- **No match:** Contoso (Russian registration) vs. Contoso GmbH

© 2026 Hawk or its affiliated companies. All rights reserved

Layer 2: Investigating What Remains

When a hit is generated, AI can also take over the investigation a senior analyst would otherwise run by hand, gathering evidence, applying policy, and documenting a clear rationale:

Payment Screening Agent – evaluates each transaction hit individually, analyzes counterparty history, and calculates a confidence score, enabling straight-through processing for high-confidence false positives.

Customer Screening Agent – runs targeted, domain-whitelisted OSINT searches (Reuters, BBC, SEC, court records) to resolve PEP and adverse media ambiguities that internal data alone can't settle.

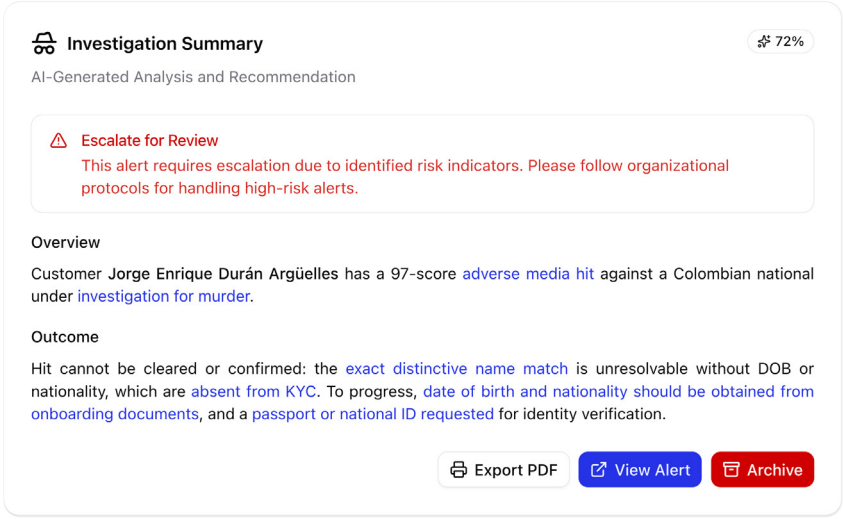

Image 2: *Example of Agentic AI Alert Review Within the Hawk Platform*

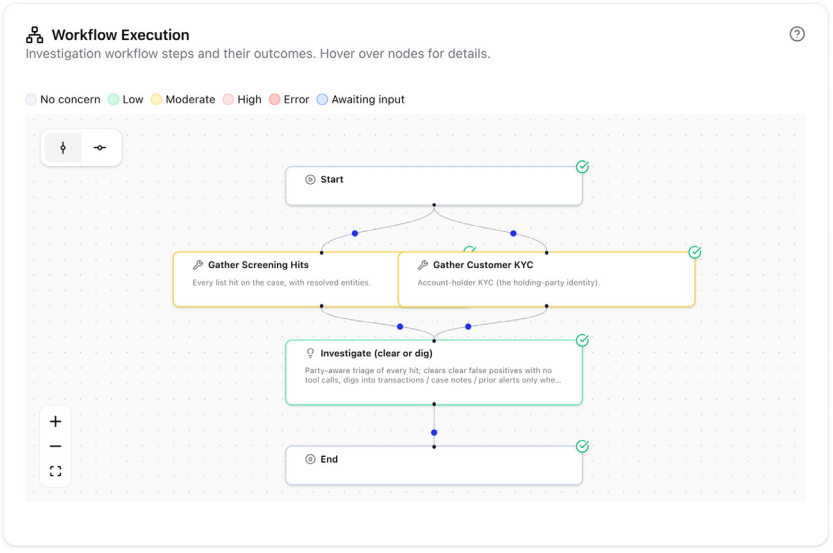

Image 3: *Example Agent Workflow Within the Hawk Platform*

© 2026 Hawk or its affiliated companies. All rights reserved

Each AI layer delivers real value on its own, but paired together, the gains compound: less noise reaches an agent to investigate, and the hits that do get resolved quicker. Table 2 summarizes how the two layers divide.

Table 2: Technical Differences Between Machine Learning vs. AI Agents

Dimension	Machine Learning at Detection	AI Agents for Alert Review
Primary Goal	Pre-block contextualization & attribute derivation	Deep multi-step hit investigation & evidence assembly
Application Domain	Payment strings & customer profile inputs	Payment hits & customer PEP/ sanctions/ adverse media cases
Execution Timing	Real-time (milliseconds)	Asynchronous/ post-hit triage (seconds to minutes)
Underlying Tech	Multilingual language models & feature engineering	LLM orchestration + whitelisted OSINT web search
Token Cost Impact	Zero token consumption (lightweight ML)	Managed token usage (triggered only on residual hits)



© 2026 Hawk or its affiliated companies. All rights reserved

Which layer to prioritize depends on where the pressure sits in your screening operation. Table 3 maps common concerns to the layer best placed to address them.

Table 3: Matching the AI Layer to Your Screening Concerns

Concern	Machine Learning	AI Agents	Best Fit
Payment delays and settlement SLAs	✓ Clears hits pre-block in milliseconds	✓ Shortens holds by resolving blocked payments in seconds to minutes	Both
Customer friction at onboarding	✓ Stops false hits before they hold a customer	✓ Resolves the remaining hits faster	Both
Missing data in payment messages	✓ Derives gender, entity type and jurisdiction from the name	–	Machine Learning
Analyst backlog and review workload	✓ Reduces how many hits are created	✓ Investigates and resolves residual hits	Both
pKYC re-screening volume	✓ Suppresses non-material background hits	✓ Reviews profiles that have genuinely changed	Both
Complex PEP and adverse media cases	–	✓ Runs targeted research and weighs the evidence	AI Agents
Audit-ready documentation	✓ Attribute-level score breakdown	✓ Full rationale, confidence score and citations	Both
Compute and token costs at volume	✓ Lightweight, no token consumption	– Token cost on every hit reviewed	Machine Learning

Each layer covers gaps the other leaves open, and combining them is where the real gains come from. Diagram 5 shows how.

Diagram 5: Why Combining AI Detection and Agents Outperforms Either Alone

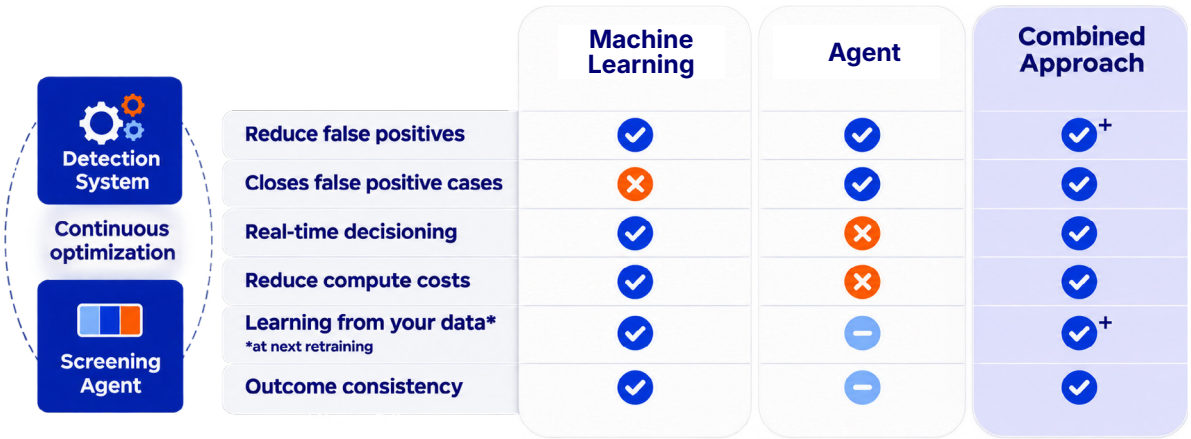

© 2026 Hawk or its affiliated companies. All rights reserved

How to Evaluate Screening AI Solutions

If you're evaluating AI for sanctions and watchlist screening, you'll quickly find that most vendors lead with the same claim: a false positive reduction somewhere in the 60–95% range. Those numbers show that AI can make an impact, but they won't help you choose between solutions. The real differences lie in the questions behind the number. How does the solution reach its decisions? Can your team explain and defend them? Will the results hold up on your institution's own data? The pillars below set out what to ask each vendor so you can tell them apart.

Pillar 1: Screening Expertise & Configurability



Screening-Specific Design: Models and agents should be built for screening, not adapted from generic AI; that means trained to handle name variants, transliteration, missing payment fields, and list-specific nuances, with logic your team can follow



Ongoing Performance Monitoring: Institutions need clear visibility into how models and agents perform over time, including hit rates, drift, and decision outcomes, so potential issues surface before they reach an audit



Self-Serve Configuration: Compliance teams should be able to adjust thresholds, rules and agent behavior directly, at a granular level (by list, segment or alert type), without waiting on the vendor for every change

© 2026 Hawk or its affiliated companies. All rights reserved

Pillar 2: Evidence & Results



Effectiveness & Efficiency: Vendors should back their claims with model-level statistics rather than a headline reduction figure, showing both effectiveness (how reliably the model keeps true matches, known as recall) and efficiency (how many false positives it removes), measured on labelled data at the confidence threshold you'd actually run, so you can confirm noise comes down without true risk being missed



Comparable Client Outcomes: Ask for case studies or benchmarks from institutions similar to yours in industry, size, region, and screening volume



Verifiable Results: Results should be backed by a clear methodology, ideally tested on your own data in a proof of concept, so you can trust the numbers before you commit

Pillar 3: Speed & Latency Criteria



Pre-Block Speed: Machine Learning must run pre-block in sub-300 milliseconds to avoid payment and customer onboarding delays



Asynchronous Agent Triage: Agents execute asynchronously post-hit, completing deep investigations in sub-minute runtimes

© 2026 Hawk or its affiliated companies. All rights reserved

HAWK

Pillar 4: Governance & Compliance Criteria



White-Box Explainability: Every agent recommendation must include step-by-step chain-of-thought logging and tamper-evident source citations



Configurable Autonomy: Configurable confidence thresholds allowing institutions to adjust STP tolerances to internal risk appetites

Pillar 5: Security & Deployment Criteria



VPC Isolation & Data Privacy: Isolated tenant execution in Virtual Private Clouds (VPC); zero PII exposure to public LLMs



Non-Intrusive Overlay: API-first overlay architecture sitting directly above existing case managers via API, MCP, or UI Connector without platform replacement



Model Choice Flexibility: Model gateway routing tasks across Claude, GPT, or fine-tuned open-source models to optimize task accuracy vs token cost

Evaluated individually, these pillars indicate whether each layer performs well in isolation. Evaluated together, they reveal something more important: whether the two layers were designed from the outset as a single, integrated architecture, or assembled afterward from separate point solutions.

Table 4: 4-Pillar Comprehensive Buyer's Best-in-Class Features

Pillar	ML Detection	Agentic Review	Enterprise Baseline
1. Speed & Latency	Sub-300ms pre-block processing	Sub-minute asynchronous triage	Zero degradation of payment SLAs
2. Governance	Deterministic model cards & drift checks	Chain-of-thought reasoning & citations	Immutable audit logging
3. Security	Native VPC execution	PII redaction gateway before LLM calls	SOC 2 Type II & ISO 27001 certified
4. Cost Control	High-volume, fixed-cost ML filtering	Task-based LLM model routing	Non-disruptive API/UI overlay

© 2026 Hawk or its affiliated companies. All rights reserved

Decision Governance: Straight-Through-Processing

The buyer's framework identifies configurable autonomy as a governance requirement; in practice, that autonomy takes the form of Straight-Through-Processing (STP) – allowing high-confidence false positive hits to be cleared automatically, without human review, once a hit's confidence score clears a threshold the institution sets itself. Raising or lowering that threshold is a deliberate trade-off between throughput and caution, which is exactly why regulators expect institutions to control, document, and defend the choice rather than treat it as a fixed setting.

How much autonomy is appropriate, and under what safeguards, differs by jurisdiction.



European Union

Filtering false positives in AML and sanctions screening generally fall outside the EU AI Act's high-risk categories, though this depends on intended purpose and on whether the system profiles individuals, which should be assessed and documented. GDPR Article 22 restricts decisions based solely on automated processing that have legal or similarly significant effects. Where an exception applies, such as authorisation under AML law, safeguards include the right to human intervention and to contest the decision.



United States

US regulators do not impose a general requirement for a human to review every screening hit. The April 2026 interagency model risk guidance (OCC Bulletin 2026-13 / SR 26-2) withdrew the 2021 statement on model risk for BSA/AML systems and explicitly places generative and agentic AI outside its scope, leaving institutions to govern AI agents through their own risk management frameworks. That puts the onus on banks to show their approach is sound: clear thresholds for straight-through processing, documented decisions, and a full audit trail remain the safest way to demonstrate it.

While the regulatory specifics vary by market, the principle of human oversight and defensibility remains.

© 2026 Hawk or its affiliated companies. All rights reserved

HAWK

Keeping Costs Predictable at Scale

Configurable governance only matters if the underlying economics hold up at scale. A primary concern with LLM-powered agents is token cost - calling an LLM against every raw fuzzy match can be expensive. This is why the hybrid architecture reserves agentic review for a reduced (~25%) population of hits machine learning couldn't resolve.

Table 5: Operational Cost & Scalability Across Review Models

Operational Model	Cost per Alert/Hit	Scalability & Token Efficiency
Manual Human Review	$5.00 – $20.00	Low scalability; severe analyst staffing bottlenecks
Agentic Review	$1.00 – $5.00	High scalability; expensive LLM token usage at volume

How To Keep Token Costs Under Control

Keeping agentic alert review affordable at scale isn't down to one setting. It comes from cost control built into how the architecture is designed, with:



Adding agents only where they add value: Machine learning resolves clear-cut false positives first, so LLM-powered agents are triggered only on the residual hits that need deeper investigation



The right model for each task: A model gateway routes each step to the most suitable model, whether Claude, GPT or a fine-tuned open-source model, balancing accuracy against token cost rather than defaulting to the most expensive option

© 2026 Hawk or its affiliated companies. All rights reserved

Hawk's Approach:
Hybrid AI That Improves Over Time

Designed from the outset as a single integrated architecture, not assembled afterward from separate point solutions, Hawk operationalizes the framework pillars through three complementary products, each covering a distinct part of the hybrid framework described throughout this paper:

01

False Positive Reduction AI (Machine Learning Model)

Operates inline, deriving gender and entity type the way this paper's earlier examples described; expected to deliver a ~25% false positive reduction.

02

Payment Screening Agent

Reviews payment hits on hold, weighing evidence against sanctions lists to release false hits in second-to-minute runtimes.

03

Customer Screening Agent

Automates complex PEP, adverse media, and sanctions investigations for customer profiles using whitelisted OSINT web retrieval.

Each layer can be adopted on its own, depending on where an institution's pressure sits, but the greatest efficiencies come from using them together. Every step removes more false positives before the next, so the hits that finally reach an analyst are the ones that need human judgment.

© 2026 Hawk or its affiliated companies. All rights reserved

Diagram 6: The Compounding Efficiency Cycle

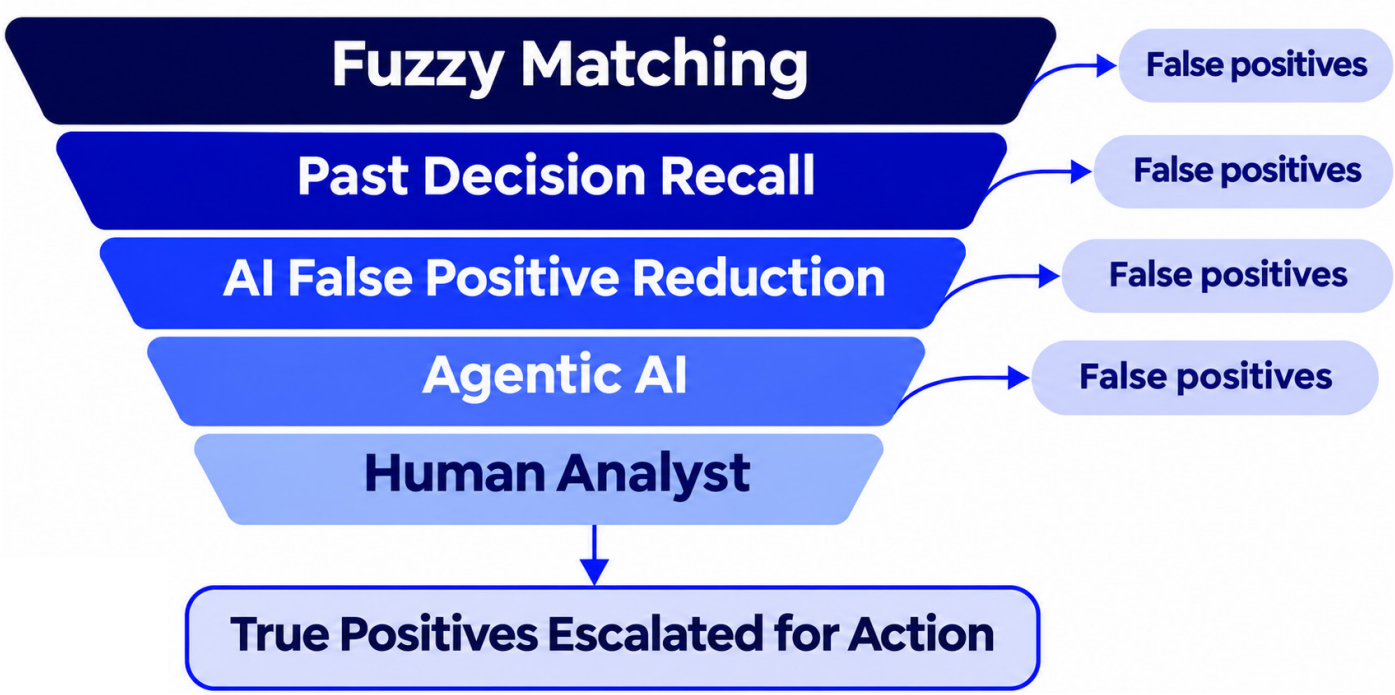

Sharpen Precision Over Time

Hawk's screening agents improve with use. Feedback on agent decisions, including where analysts confirm or override a recommendation, is reviewed and improvements can be applied at the next scheduled review, so agent performance keeps improving and the effect compounds review after review. Over time, more hits can be resolved with confidence, freeing analyst capacity for the harder cases and bringing cost per alert down. Every change stays inside your existing scheduled review and change-control process, with a full audit trail.

Looking further ahead, the same principle could extend upstream. Outcomes from agent and analyst review could inform how fuzzy matching itself is tuned, so fewer low-value hits are generated in the first place. Rather than only resolving false positives faster, the screening system would learn to create fewer of them.

© 2026 Hawk or its affiliated companies. All rights reserved

Real-World Impact

Two independent deployments demonstrate this architecture operating under real production conditions:

False Positive Reduction AI

Benchmarked against a European bank's incumbent provider, Hawk's model caught 8 percentage points more sanctions risk while generating 27 percentage points less noise.

Payment Screening Agent

Alviere, a US-based embedded finance platform that lets retail brands offer banking, card, and payment products through bank partnerships, works with enterprise customers across a broad range of markets and customer segments. As its enterprise customer base and transaction volumes grew, Alviere broadened its screening rules to avoid missing genuine risk, but broader screening meant far more alerts entering the review queue, the majority of them false positives still requiring manual review before a case could be closed. To reduce that manual effort, Alviere tested Hawk's Payment Screening Agent head-to-head against an alternative agentic AI provider on a set of recent counterparty screening alerts. Hawk's agent confidently and correctly cleared 96% of alerts, and for the remaining 4%, correctly recognized the evidence was insufficient and escalated for human review – catching several cases the alternative provider missed, while also providing a more comprehensive disposition narrative for Alviere's compliance team. Results were measured within two weeks of deployment.

Read more in the case study on hawk.ai.



© 2026 Hawk or its affiliated companies. All rights reserved

Built to Fit Institutions of Any Size

Hawk's capabilities serve any bank or payment provider running screening today. The False Positive Reduction model delivers immediate ROI regardless of transaction volume; the AI agents deliver the greatest capacity gains for operations teams of three or more analysts, while remaining fully deployable for smaller teams.

Explore It Yourself

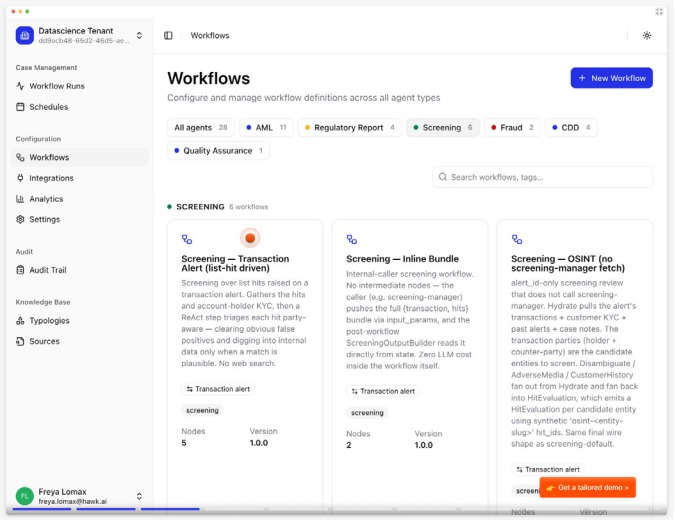

Interactive Demo 1:
Payment Screening Agent Walkthrough

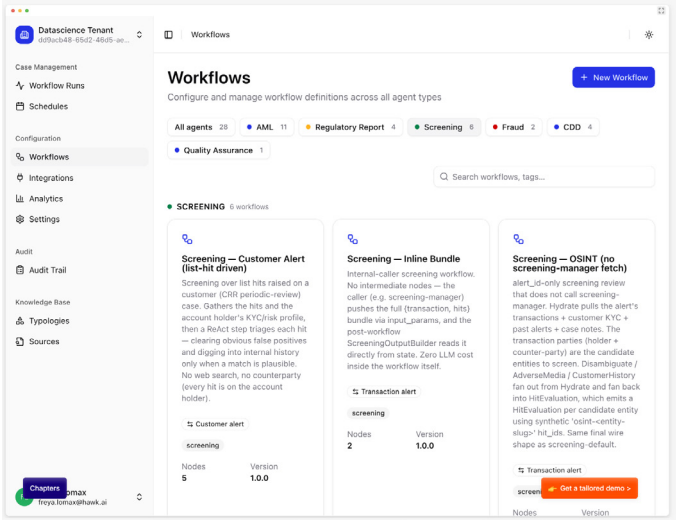

Interactive Demo 2:
Customer Screening Agent Walkthrough

© 2026 Hawk or its affiliated companies. All rights reserved

Future Outlook

The governance and cost frameworks outlined in this whitepaper describe where the hybrid architecture stands today. With 84% of institutions already planning to increase agentic AI investment over the next two to three years, the next layer of this architecture is already coming into focus.

Quality Assurance (QA) Control Agents, Multi-Agent Collaboration & Unified CDD

As agents take on more of the decisioning at scale, a new oversight need follows close behind: who, or what, reviews the agents themselves. Quality Assurance Control Agents extend the same investigative pattern described throughout this paper, but turn it inward, running continuous, autonomous checks over both human and agent decisions to catch policy drift or emerging bias before it compounds, rather than relying on periodic manual sampling.

The second extension is structural rather than supervisory: connecting payment screening, customer screening, Customer Due Diligence (CDD), and perpetual risk scoring into a single, continuously updating view of customer risk, instead of four systems that each see only their own slice of it. The pKYC suppression and escalation logic described in the hybrid AI framework is a preview of this shift – the same discipline, applied across every system that touches a customer's risk profile rather than one at a time.

© 2026 Hawk or its affiliated companies. All rights reserved

Summary

Screening teams today have outgrown the systems built to support them. Sanctions lists keep expanding, liability keeps tightening, and the shift to continuous pKYC re-screening keeps adding volume; all while fuzzy matching, still the backbone of most screening programs, continues to generate the same 90–95% false positive rate it always has.

Pairing machine learning with AI agents changes that equation: non-material hits are suppressed before they're created, and the hits that remain are investigated with a depth at speed, all while preserving the audit trail regulators expect. For compliance leaders weighing where to focus investment, that combination – lower operating cost, less analyst burnout, and a stronger risk posture – doesn't have to be a trade-off.

To learn more, contact us at hawk.ai to book a discovery session or ROI workshop to explore the potential benefits of Hawk's hybrid AI screening solution for your institution.

About Hawk

Hawk is the leading provider of AI-supported anti-money laundering, screening, and fraud prevention technology. Banks and payment providers globally use Hawk's modular platform to pinpoint financial crime risk with precision, cut operational overhead, and ensure regulatory compliance.



© 2026 Hawk or its affiliated companies. All rights reserved

HAWK

Hawk AI GmbH
Friedenstrasse 22B/i3
81671 Munich
Germany

Hawk AI USA Inc
230 Park Ave, Floors 3 & 4
New York, NY 10169
U.S.A.

Hawk AI APAC Pte Ltd
160 Robinson Road, #14-04
Singapore Business Federation Center
068914 Singapore

Hawk AI Technology UK Ltd
Birchin Court, 5th Floor 19-25 Birchin Lane
London
EC3V 9DU
United Kingdom

Follow us on LinkedIn >

info@hawk.ai >

www.hawk.ai >