Shufti

The Context Gap in AML Risk Assessment

A financial crime practitioner view of
AML screening and ongoing monitoring.

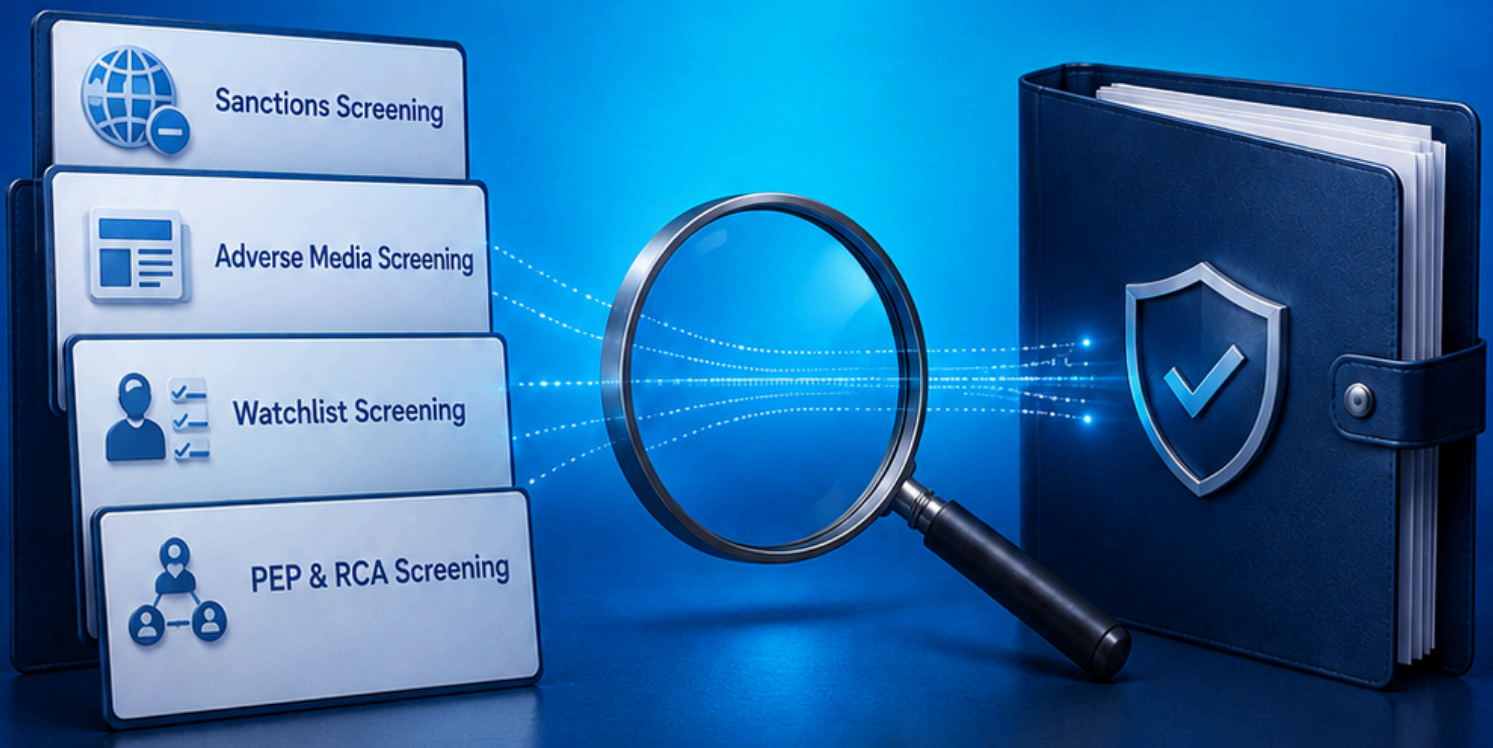

FINDINGS FROM SHUFTI'S VOICE OF CUSTOMER RESEARCH · 600+ ORGANISATIONS

KJM Age Verification · PCI DSS · iBeta Level 3 · ISO/IEC 27001:2022 · CE+ · QG GDPR · SOC2 · GDPR · CCPA

In this report

The real challenge of solving detection and decision

AML-obligated organisations worldwide continue to face rising compliance costs despite sustained investment in compliance solutions. While these investments have strengthened sanctions and PEP screening, many solutions still struggle to support risk assessments aligned with an institution's own risk exposure, risk appetite and jurisdictional obligations.

To understand why, Shufti's proprietary Voice of Customer research analysed feedback from over 600 organisations evaluating or operating AML programmes. The findings show that **51% identified recurring manual handling during ongoing monitoring as a key operational challenge**, making it the most frequently cited control-specific issue in the study.

Risk screening capability has matured considerably. Despite this, the challenge of reaching the defensible decision has not been addressed yet.

Every AML programme seeks to establish customer identity, assess risk, and maintain that understanding throughout the relationship. Yet most screening systems answer only a narrow question at a single point in time. The verified identity established at onboarding rarely flows into screening with context, same risk profile may not show in secluded transaction monitoring system, leading to false alerts in ongoing monitoring both in customer risk and transaction risk. As a result, compliance officers have to spend a lot of time in dissecting the right decisions through risks.

This pattern appeared consistently across the research. While **32% identified manual handling as the single largest operational challenge**, the same underlying issue surfaced across sanctions resolution, PEP verification, ongoing monitoring, and transaction monitoring. Four separate controls pointing to the same condition indicate an architectural limitation, not isolated process inefficiencies.

Shufti

Voice of AML compliance professionals

The same challenge of manual handling of alerts surfaced independently in different controls, attached each time to a different task.

Sanctions screening, PEP screening, ongoing monitoring, and transaction monitoring were each raised as separate operational areas, and each returned the same condition attached to a different task. The figures below are read together rather than in sequence, because the recurrence is the finding.

32%	ACROSS THE PROGRAMME	Named **manual handling** their single largest operational difficulty, before any individual control was examined.
45%	SANCTIONS SCREENING	**Resolution.** Clearing an individual match, one alert at a time, because deciding on those alerts takes a lot of time as context and relevance are misisng.
50%	PEP SCREENING	**Verification.** Establishing risk manually because difficulty in deciding on common name alerts in ongoing monitoring of PEPs.
51%	ONGOING MONITORING	**Recurrence.** Re-screening on a schedule as a substitute for continuous monitoring, because too many alerts becomes difficult to handle manually for compliance offficers
20%	TRANSACTION MONITORING	**Reconciliation.** Assembling three vendors' separate assessments of the same customer into one risk view.

How manual AML controls compound across the customer lifecycle

Incorrect risk assessment at onboarding due to false positives overload can lead to higher alerts in ongoing monitoring requiring even more manual effort

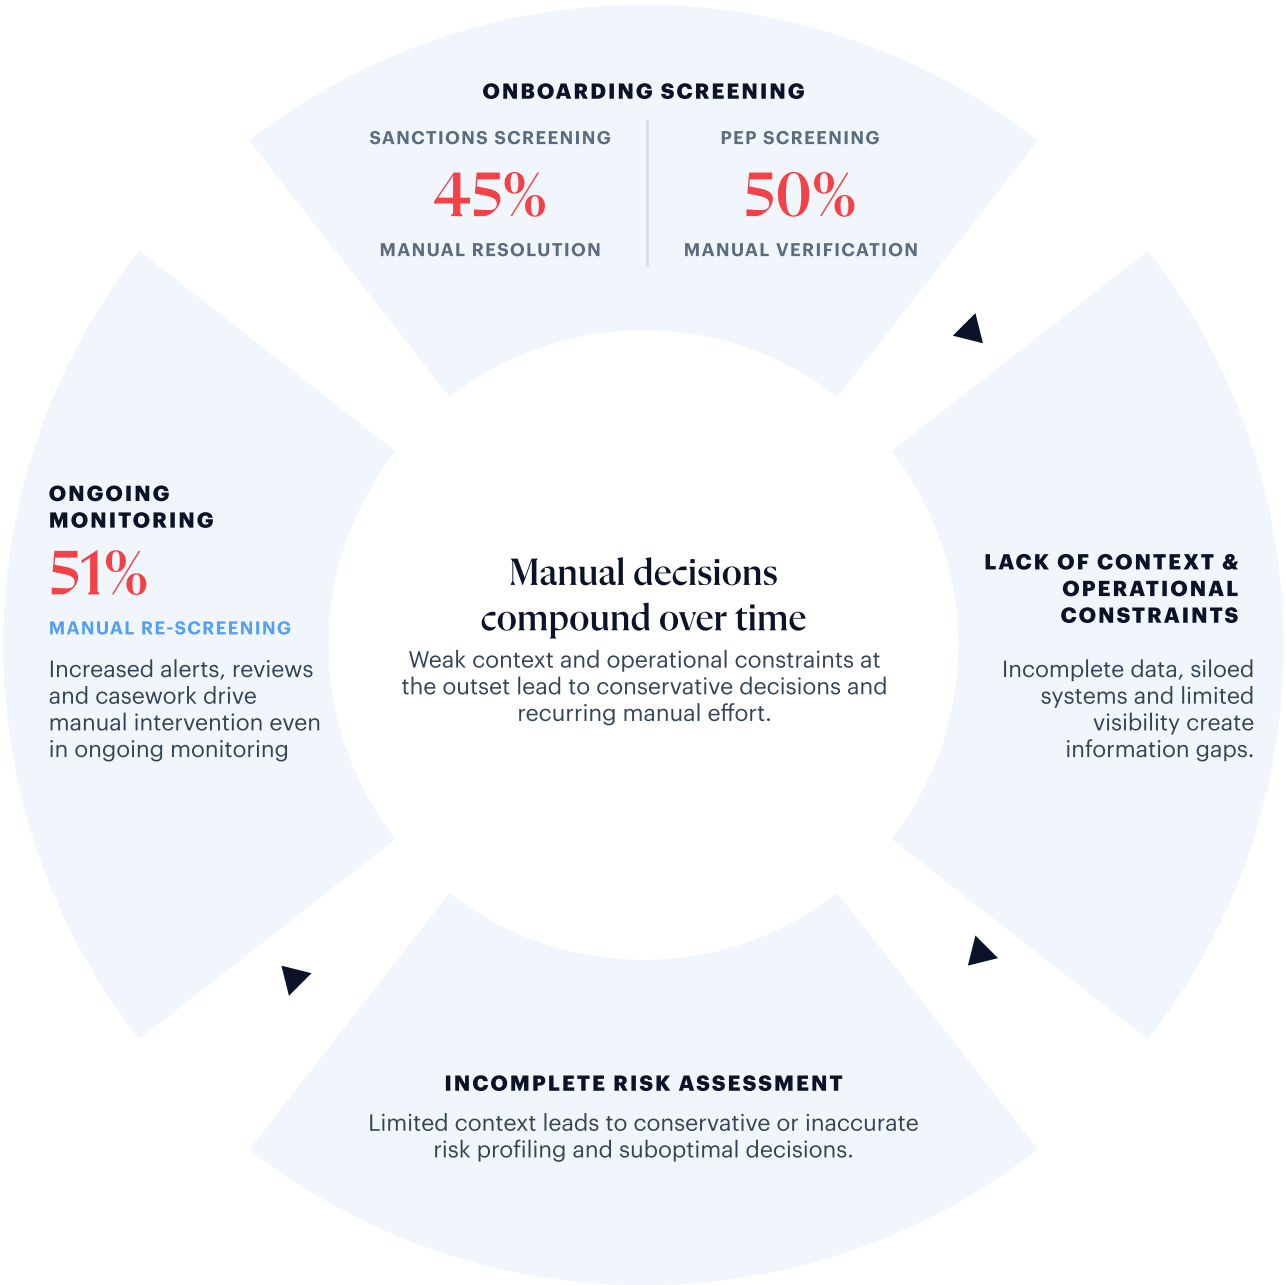

Figure 1. Weak context at the outset compounds clockwise through the lifecycle. Percentages are coded from the dataset.

Where the findings come from

This report is based on an analysis of anonymised market intelligence derived from Shufti's customer and prospect base, comprising more than 600 organisations across the globe actively evaluating anti-financial crime solutions, including identity verification, fraud prevention, AML compliance, and transaction monitoring technologies.

The analysis draws on insights shared by fraud, risk, compliance, and financial crime practitioners during solution evaluation and commercial discovery engagements. These discussions reflect real-world operational challenges, technology priorities, and evaluation criteria raised in the context of active business initiatives. All responses were anonymised prior to analysis and reviewed using thematic analysis to identify recurring patterns and emerging trends across industries, regions, and organisation types. Findings are presented in aggregate to provide a broader view of market priorities while ensuring that no individual organisation or respondent can be identified.

600+
ORGANISATIONS ANALYSED

Global
INDUSTRIES, REGIONS, ORG TYPES

Thematic
ANALYSIS, MULTI-CODED

HOW TO READ THE PERCENTAGES

Percentages represent the proportion of organisations raising each challenge or issue during these engagements. Organisations frequently raised multiple operational challenges during their solution evaluations, and many themes emerged independently from the same conversations. The themes presented are multi-coded and the percentages are not intended to sum to 100%. This approach captures the breadth of operational concerns practitioners face rather than forcing a false choice between competing problems.

Regulators are now measuring the cost of compliance itself

Regulators themselves now recognise that mounting regulatory pressure may be driving AML compliance costs higher without a matching gain in risk reduction.

FinCEN AML/CFT Compliance Survey

In September 2025 the United States Financial Crimes Enforcement Network (FinCEN) opened a formal Survey of the Costs of AML/CFT Compliance, seeking data from financial institutions on direct compliance costs, labour, monitoring systems, vendor fees, and the share of operating expenditure AML compliance represents. The survey results are not yet published, but the industry responses submitted already name personnel and technology as the largest components of compliance cost, and manual workflows and multi-vendor integration among the principal operational burdens.

Federal Reserve community banking study

~10%

A Federal Reserve Bank of St. Louis study of 1,091 community banks found total regulatory compliance costs near **10% of non-interest expense at the smallest banks**, roughly double the largest, with the Bank Secrecy Act the most burdensome area at about a fifth of that total.

Independent industry benchmarking

As per PwC's EMEA AML Survey 2026 more than five hundred financial institutions across Europe, the Middle East, and Africa found AML compliance costs rising structurally: up ten to thirty percent over two years for a third of institutions, customer due diligence now the largest bottleneck and judged too rules-based by more than forty percent of respondents, confidence in transaction monitoring roughly halved since 2024, and data quality the leading barrier to AI adoption at up to 89%. Those figures measure the problem from the outside; this report's findings explain it from the inside.

Why AML compliance becomes costly and inefficient

AML compliance becomes costly and inefficient because of a structural mismatch between how customer data flows and how risk decisions are made. The customer's identity is verified once, at onboarding, and the information captured there, date of birth, nationality, document details, associated risk status, is exactly what resolves a screening match. Yet when an alert arrives, that information is not there: collected, verified, and paid for, but behind a different vendor's interface, so the officer reconstructs context by hand. Staffing changes how fast the work is absorbed, not whether it exists.

The cause is architectural.
Detection and decision therefore operate as separate functions. The system detects with growing accuracy, but a detected alert cannot be resolved without manual context gathering, so more alerts mean more manual work and more headcount. The cost lands in four connected places, all with the same root cause.

WHERE THE COST LANDS	WHAT HAPPENS
Decision time	The information needed to resolve an alert is either missing or manual overview may take a lot of time. Officers search for dates of birth, nationality, document numbers, and transaction history, spending a lot of time in investigating the risk relevance.
Analyst hours and headcount	Context gathering scales with alert volume. When volume outruns the team, cost increases by hiring, so spending scales with noise rather than actual risk.
Backlog and delay	Unresolved alerts pile up. Approvals and clearance wait for manual resolution of alerts that would close automatically if context were preserved.
Regulatory risk	Under backlog pressure, teams clear faster, which weakens the control. Decisions made on fragmented context are inconsistent, undocumented, and unapproved.

Findings

Across AML operations, the resolution of alerts was raised as a difficulty more often than their detection. Compliance officers share the challenge of spending time manually gathering context because the information required to resolve a screening match is not present when the alert arrives. That context, verified identity data from onboarding, transaction history, previous risk assessments, exists within the organisation's systems but does not flow to the point where decisions must be made.

The four control-level findings that follow each describe manual work, and that is deliberate rather than repetitive. What is manual differs in every case. Read in order, they show one architectural failure expressing itself as four operational problems, and the two findings that close the chapter explain why.

45%
experience friction in sanctions screening resolution

50%
of PEP screening practitioners report manual verification challenges

51%
experience friction in transaction monitoring reconciliation

21%
lack sufficient context during alert resolution

46%
lack sufficient context during alert resolution

27%
face increasing regulatory pressure

Sanctions screening

45% of organisations face challenges because context gathering for match resolution as slow and manual

For sanctions, coverage gaps were raised by 18% of organisations while the manual resolution of matches was raised by 45%. The difficulty named most often is therefore the need to gather context to resolve a match, because the alert itself lacks the information required to decide whether this customer is the individual named.

THE MANUAL TASK · RESOLUTION

An alert already exists and a decision must be reached on it. The officer is **adjudicating a candidate match** that the system raised and could not close.

Clearing matches is manual and slow — 45%

Officers work through alerts one by one, checking matched names for unique identifiers and additional context which may be beyond capability of their existing solution.

Coverage gaps, local and regional lists — 18%

Close to a fifth of organisations still report gaps at local and regional level. This implies that businesses may not be able to do sanctions screening as per their AML risk exposure.

Cannot distinguish true from false match — 11%

This is fundamentally a lack of context problem arising due to common names, names of sanctioned entities matching with irrelevant lists and missing primary and secondary exposure lablels.

PEP screening

50% of organisations face challenges due to checks as manual or ad hoc because decision context is missing

For PEP screening, officers resolve checks case by case, resolving the names similarities with non-PEP customers.

THE MANUAL TASK · VERIFICATION

Verification rather than resolution. Sanctions work asks whether a specific alert applies to this customer. PEP work asks a prior question, **whether political exposure exists at all**, and answering it sends the officer outside the system entirely, to open sources and web search, because no verified record inside the platform can settle it.

Checks are manual or ad hoc **50%**

The system cannot close the check automatically because it lacks the verified identity data needed to resolve the match.

Coverage, accuracy or consistency gaps **33%**

This spans two distinct issues, lists not held and lists whose contents cannot be relied upon. The distinction matters operationally, because the second persists even where coverage is complete.

Ongoing monitoring

51% of organisations face challenges in resolving alerts

Ongoing monitoring is where fragmentation extends across time rather than sitting within a single decision. A customer's risk profile, established at onboarding, must be maintained continuously throughout the relationship. What is meant to run continuously is being performed manually, at intervals, because of unnecessary alerts due to poor quality of customers risk data.

THE MANUAL TASK · RECURRENCE

The two findings above concern a single event, an alert that must be cleared or a status that must be confirmed. This one concerns time instead. **Nothing watches the customer between scheduled runs**, so a person must decide when to look again, and the control degrades from continuous to periodic while the risk may change continuously.

Manual workaround instead of automation (51%). Practitioners must re-screen customers on a schedule rather than monitoring activity in real time.

Identity and transaction history not connected

The system that verified identity at onboarding is disconnected from the system monitoring transactions. Customer context cannot flow from one to the other.

Risk profile not maintained continuously

The risk profile must stay current for the duration of the relationship, but it remains static because data is not updated and this happens often because vendor may not have control over the data layer.

Previous work not remembered across cycles

Each screening run treats the customer as a new case with no memory of previous work. Alerts resolved before are investigated again, wasting analyst time.

Beneficial ownership adds structural complexity

20% cite structural complexity as the principal difficulty. Ownership changes less often than transactions but is harder to verify, and re-verification flows into a separate system.

Transaction monitoring

20% of organisations face challenges due to separate vendors as their principal operational difficulty

Transaction monitoring is where fragmentation between suppliers becomes most visible, raised by 20% of organisations as their principal operational difficulty. Officers must manually reconcile three vendors' risk assessments of the same customer.

THE MANUAL TASK · RECONCILIATION

The preceding three findings describe manual work inside one system. This one describes manual work between systems: the officer is **assembling a single view of one customer from suppliers** that each hold part of the answer and cannot see the others.

No vendor consolidation

An officer running monitoring, sanctions, and identity verification at three vendors must reconcile three separate risk assessments to answer a single question.

Disconnection from customer risk profile

Monitoring has no visibility into the customer's risk profile as defined in AML Risk exposure and processes against generic rules, leading to false alerts. This is likely to happen because alerts are not backed by regulatory logic.

Volume processed without risk-aligned logic

Without context based on regulatory logic, every alert is equally suspicious and equally slow to resolve. More volume triggers more alerts and more manual decisions.

Gaps across payment channels

Most solutions screen conventional payments well but lack visibility into cryptocurrency, digital wallets, or payment apps, leaving risk assessments incomplete.

Jurisdiction-specific rules cannot be embedded

Risk varies by jurisdiction, but systems apply generic rules everywhere. Officers face a false choice: relax rules globally or tighten and create alert overload.

Real-time performance lags transaction flow

Alerts arrive after transactions move or rules run overnight on batch, so suspicious transactions may clear before detection.

46% of lack of context

driven false positives are due to common name noise leading to poor decisioning

The most direct evidence of context fragmentation is the absence of critical information when an alert is raised. Practitioners more often reported that clearing alerts was slow and manual than that data was missing. One reading of that pattern is that missing context surfaces as missing time, in hours spent searching for information the organisation already possesses.

Common name ambiguity without verified identity data

A common name is ambiguous only where insufficient data exists to distinguish individuals. The verified date of birth, nationality, and document identifiers captured at onboarding would resolve these ambiguities instantly, yet they do not arrive with the alert because they sit in a different system.

Verified identifiers not flowing to the screening system

The verified date of birth, nationality, document number, and identity source are collected, verified, and paid for at onboarding. None of this reaches the screening system when the alert arrives, forcing manual reconstruction of context.

Manual searching for information that already exists

Officers search for dates of birth, nationality, document numbers, and transaction history: data that exists within the organisation's systems but is not connected. A team whose alerts arrive without a date of birth does not report an identifier problem. They report that matching takes longer than it should, and that time is spent searching for context.

Lack of context leads directly to **longer decisioning time** , since every alert must wait while an officer reconstructs what the organisation already knows. That added time compounds into **inefficiency**, as analysts absorb work the architecture should have carried. And inefficiency, sustained across thousands of alerts, becomes **higher cost ,** paid in headcount that scales with noise rather than risk.

Regulatory obligation

27% of organisations cite regulatory and audit pressure as their principal challenge

Regulatory expectation and customer volume are rising together, while the control absorbing both is an analyst reviewing alerts individually. The convergence creates multiple simultaneous pressures on a single bottleneck that cannot scale.

Regulatory and audit pressure — 27%

Regulators are asking for more: more screening lists, more frequent reassessment, more documentation of due diligence. This obligation is growing faster than the systems or staffing can accommodate.

Growth and customer volume pressure — 19%

As customer portfolios grow, alert volume grows with them. The team absorbing both regulatory pressure and volume growth is the same analyst reviewing alerts individually, the one component that cannot scale.

Speed and friction pressure — 18%

Compliance requires quick decision-making on alerts, yet manual context gathering makes speed impossible. The friction from context fragmentation compounds the pressure to clear alerts faster.

Backlog blocking approvals and oversight — 8.5%

When the team cannot keep pace, alerts pile up unreviewed. Every hour of backlog is exposure to a risk already detected, making it a regulatory and operational risk position.

Why AML screening becomes the manual work

1

The data going into the match is weak

A record built from a verified document carries the name as printed, a date of birth, a nationality, and machine-readable data; one built from a sign-up form carries whatever the customer typed. A strong identity verification layer which reconciles names and establishes the right name match is pre-requisite.

2

The matching logic is a blunt trade-off

Matching narrowly misses the genuine hit spelled differently; matching loosely returns a flood. The solution should have a high quality data to resolve a common name match and ensure accurate risk.

3

The alert arrives without the context to resolve it

The system hands the officer a possible match, not an answer. Deciding needs dates of birth, nationality, document number, and listing source and most important updated PEP, Sanctions and Watchlist databases for higher accuracy.

4

The systems do not share what they know

Identity, screening, transaction monitoring, and ownership checks typically sit with four suppliers, none sharing a record. The verified date of birth and document number captured at onboarding may not always reconcile with screening system calling for vendor consolidation.

5

The unnecessary Ongoing Monitoring alerts due to poor data quality

The ongoing monitoring procuces a lot of alerts if the underlying data doesn't reflect the actual risk, this is likely to compound even with AI power alert triage if the data under the Agentic AI layer is not informed enough.

AML screening needs a decisioning layer

Most AML architectures carry two layers. A data layer holds the lists, the verified identity records, and the transaction history. An application layer matches names, applies rules, and routes what it finds. A decade of investment has gone into both, and both have advanced. What sits between a detection and a decision is a third layer that most architectures do not have, so converting one into the other falls to a person. It is why **51% of organisations named recurring manual handling in ongoing monitoring as an operational challenge**, the most frequently cited control-specific difficulty in this research.

What an AML Decisioning Layer Requires

Context drawn from the verified identity record, so that an alert arrives carrying the identifiers that resolve it. Data that stays current, with sanctions and PEP changes applied as they occur rather than at the next scheduled cycle. And logic that moves on two axes at once, tracking regulatory change while remaining configurable to the institution's own risk exposure and risk appetite. Where those hold, context can be assessed as the alert is raised, and most alerts resolve without manual work, leaving people the cases that call for judgement.

The question that matters most

Which lists a programme screens against is a necessary question. A second carries more weight, which is whether the system that establishes who the customer is and the system that decides what that name signifies are the same system. Where they are, context is preserved across the whole AML lifecycle, from customer risk assessment through transaction screening to ongoing monitoring, and that is what makes an assessment accurate, explainable, and aligned with the institution's own risk appetite. Most organisations have never been asked it, and the checklist that follows exists so that they can ask it of themselves.

One system, one record, one AI Reviewer, the whole lifecycle

For most suppliers the system that verifies a customer's identity and the system that screens and monitors that customer are not the same system: the platform is assembled from parts, and every join is a point where context fails to cross.

Shufti owns the full stack from identity verification through AML screening and ongoing monitoring, enabling four connected capabilities the assembled-from-parts model cannot deliver. Risk assessments are based on the data which updates continuously in response to changing regulations and updates as per a business's AML risk exposure as identified in their AML Compliance Program. The approach leads to relevant and quickly resolvable alerts.

1. **Identity read at source**

   Verification is built to each country's own documents and regulatory logic, across 240+ countries and territories with OCR trained natively on 150+ languages. A name is read in its own script, so the identifiers that later resolve a match are established at source.

2. **Screening over the same verified record**

   Sanctions, PEP, adverse media, and watchlist screening run over the verified identity record, so the context captured at onboarding is present when the match is made, and the regulatory logic is maintained in step with regulatory change rather than waiting on a third party.

3. **Monitoring that keeps the record current**

   Identity, screening, and monitoring share one record across User AML Screening, Business AML Screening, and Transaction Trust Screening. A flagged transaction updates the customer's risk profile immediately rather than waiting for the next scheduled re-screen: the continuous linkage regulators expect, available as cloud, Local Cloud, on-premises, or hybrid through a single integration.

4. **The compliance team's first AI Reviewer**

   Shufti's MLRO AI Agent triages alerts over the same verified record, delivering pre-scored evidence so analysts review decisions rather than raw matches, with judgements that stay explainable and audit-ready.

AML Vendor Assessment for Faster Alert Resolution

Each question surfaces a specific operational weakness found common in deployed systems.

ON ALIGNMENT TO RISK EXPOSURE AND APPETITE

- [ ] Can thresholds be configured and evidenced by risk tier, rather than to a fixed vendor default?
- [ ] Can screening sensitivity be tuned without a vendor change request, with every change logged for audit?
- [ ] Can the solution apply different screening rules by jurisdiction, product line, or customer segment?

ON DATA QUALITY ENTERING SCREENING

- [ ] Are identifiers (date of birth, nationality, document number) drawn from the same platform that verified the customer's identity?
- [ ] Is a name matched in its original script, or only after conversion to Latin characters, and which romanisation standard is applied?
- [ ] Where a name reaches the system already transliterated, misspelled, or romanised to a different standard, does matching still resolve to the correct entity, and can the vendor evidence that in a live test using your own records?

On context delivered with alerts

THE IDENTIFIERS BEHIND THE MATCH

- [ ] Does the alert deliver the matched identifiers in full, including the name as recorded on the list and as verified for the customer, date of birth, nationality and country of origin, place of birth, document type and number, and known aliases, and does it say so explicitly where the list holds no value rather than returning an empty field?
- [ ] Are those identifiers presented alongside the customer's own verified equivalents, with a rationale showing which fields matched and at what threshold?

SANCTIONS CONTEXT

- [ ] Does the alert identify the regime, the issuing authority, the listing instrument, and the designation date, and show whether the designation binds the institution in its own jurisdiction?
- [ ] Are the measures themselves labelled, distinguishing an asset freeze from a primary sanction, a secondary sanction, a sectoral restriction, an arms embargo, a travel ban, or an export control designation?
- [ ] Where the match is reached through ownership or control rather than direct designation, does the alert state the ownership percentage and the path to the designated party?

PEP CONTEXT

- [ ] Does the alert state whether the individual is a domestic PEP, a foreign PEP, or a PEP within an international organisation, together with the position held, the issuing country, and whether that position is currently or formerly held?
- [ ] Where the match is to a family member or known close associate, does the alert name the principal and the nature of the relationship?

OVERALL

- [ ] What proportion of alerts can be resolved from the information delivered in the alert itself?

Explainability, consolidation, continuity

ON EXPLAINABILITY AND CONSOLIDATION

- [ ] Can the solution produce a documented rationale for each decision showing why a match was cleared or escalated?
- [ ] Are screening models owned by the vendor or licensed from a third party?
- [ ] How many separate suppliers underlie identity, screening, monitoring, and corporate-structure functions?

ON CONTINUITY BETWEEN SCREENING AND TRANSACTIONS

- [ ] Does the solution assess identity screening and transaction activity against a single, shared understanding of customer risk?
- [ ] When a customer newly appears on a list or behaviour changes, does the system automatically re-evaluate risk?
- [ ] Can the solution evidence ongoing monitoring is taking place in the form a regulatory audit requires?

ON AI-BASED ALERT TRIAGE

- [ ] Does the solution have an Agentic AI based alert triage feature?
- [ ] Does the solution own the AML Risk Data Layer?
- [ ] Does the solution immediately update the risk data after a change in sanctions and PEP status falling under risk exposure lists?