

Your Edge Cases Are Our Default

Giving product teams the agility to adapt to compliance requirements without compromising conversion.





Do compliance teams and product owners have the tools and flexibility to implement that risk-based approach quickly as regulations, markets, and fraud patterns evolve?

“

Financial institutions should identify, understand, and assess their financial crime risks, and apply controls that are proportionate to those risks — prioritising effectiveness over a one-size-fits-all, rules-based approach.

— Adapted from the **Wolfsberg Group** Guidance on the Risk-Based Approach (2026)

Adapting identity verification to continuous regulatory change

How the right KYC solution lets product owners implement new compliance requirements quickly, without losing genuine users. Meeting the goals of compliance and conversion at the same time.

The challenge for compliance and product teams

Regulatory change is continuous. While compliance teams decide the new AML program, product teams are the ones who implement it, and it is the architecture of their KYC solution which decides whether each change is implemented quickly or takes time — impacting compliance and customer experience in new markets and for new business lines. The mechanics are simple:

Regulations change, and so do the compliance team's requirements

Sanctions lists are revised, markets are added, regulators raise thresholds, and new product lines shift risk. Each change has to be reflected in the KYC solution, so that compliance does not come at the cost of conversion.

Product teams face the burden of implementing this change

In most organisations the work of implementing it falls on the KYC product teams, who have to ensure the technology is readily customised to assess risks according to new rules.

A KYC provider that outsources creates roadmap friction for you

A provider that does not own its technology relies on third-party engines behind the scenes. Neither you nor the provider can alter the underlying logic — a new script, a revised sanctions rule, or a new document type becomes a request passed to an upstream vendor, and a wait on that vendor's roadmap.

An owned stack turns change into configuration

When the KYC provider owns both the identity verification and the AML stack, the same change becomes a configuration or a scoped engineering task delivered by the party that built the system. Most regulatory changes stop being engineering work.

This distinction has measurable consequences for launch speed, engineering load, conversion, and compliance risk. The sections that follow examine the challenges KYC product owners face at digital-first businesses — across crypto exchanges, remittance companies, money service businesses, fintechs, digital banks, iGaming, and mobility platforms — from name matching and transliteration to Proof of Address normalisation, compliance geofencing and sanctions screening, and data residency. For each, they show how much time the challenge consumes today, and how most of these problems are, in the end, a matter of choosing the right KYC solution.

Identity verification as a continuous reconciliation challenge

For a compliance team, identity verification can look like a settled requirement: verify a genuine document, reject a fraudulent one. For the product team, the work is far more pronounced. The harder job is reconciliation — matching names against databases, which is tougher when transliterating non-Latin to Latin scripts. This matters more than it first appears, because in many scenarios failure to reconcile leads to rejecting genuine customers and accepting false ones.

Distinguishing a forged document from a genuine one is a necessary part of identity verification, but it is not the part that determines whether a legitimate customer is accepted. That depends on reconciliation: aligning the identity a person presents with the records held about them across banks, government registries, sanctions lists, and address databases. Product teams need to ask if their KYC provider is really capable of addressing this challenge.

Two failure directions, both costly

GENUINE, REJECTED

A document can belong to a legitimate owner and still fail verification, because the name it carries does not match the name the customer entered or the record held about them. If that name is not reconciled, a genuine customer is rejected.

PROHIBITED, ACCEPTED

A document can be entirely genuine while the person presenting it belongs to a sanctioned jurisdiction. A sound solution must establish not only that the document is real, but that the issuing country is not sanctioned under the applicable AML programme. If it does not, a prohibited user is falsely accepted.

Why regulatory change becomes the product team's problem

Compliance forwards a regulatory update that has to be reflected in onboarding before the next release. Your engineers are mid-sprint on revenue features, the timeline is not yours to set, and none of the change is optional. At that moment the only question that matters is whether your verification platform lets you make the change yourself, or makes you wait because of a third-party dependency.

Why this matters

Compliance writes the rule, but product teams ship the features that ensure compliance happens according to those rules. When a requirement moves and your platform cannot be customised, your roadmap slips and your pass rates drop while you wait on a third party.

The challenge for the product team

Name matching is not something a platform simply supports or does not support for a market. It generally works at onboarding, and the difficulty surfaces later, in production, when a naming convention such as Kurdish transliteration starts producing mismatches that reject genuine customers. What matters then is whether the platform can adjust its name-matching logic on demand and implement the change your product team specifies, when you specify it. The question is not how many features a platform lists — features are overtaken as requirements change — it is how quickly the platform can be adapted to what your product now needs.

EXHIBIT 1 · WHAT AN ADAPTABLE VERIFICATION STACK DELIVERS

CAPABILITY	OUTCOME FOR THE PRODUCT ORGANISATION
Adapts matching, address, geo and deployment logic without a vendor release	Faster market launches, higher pass rates and more conversions
Requirement handled by the provider or by configuration	Fewer engineering dependencies
Change delivered as a scoped task	Shorter implementation cycles
Rule changes absorbed without re-planning	Less roadmap disruption
Fewer false rejections of genuine users	Higher customer conversion
Changes reflected without a queue	Lower risk due to non-compliance

Transliteration of non-Latin identity documents and difficulty in name matching

The trigger for your team. Your business is expanding into the MENA region, and compliance reports that Arabic and Kurdish name variations are driving sanctions-screening hits and name-matching failures. Genuine customers are being rejected while screening returns noise. This is not a policy decision. It is a change to how names are read and matched, and it has to be in place before launch.

Why this matters

Name mismatches are among the most common verification failures, and they are rarely fraud. They quietly reject genuine customers and break sanctions screening at the same time. When a new script or naming convention starts failing in production, you need the matching logic changed on your timeline, not the vendor's — which is what happens if the solution relies on a third party.

What changes, and why it hurts twice

Names frequently fail to reconcile across borders because the same individual is recorded in different scripts and conventions in different systems. It is among the most common verification failures, and it is rarely associated with fraud.

The consequence for AML compliance is direct. Sanctions, PEP, and watchlist screening operate by matching a customer name against list entries, and those lists carry their own transliterated spellings. If a name cannot be resolved to a consistent form, screening it against a sanctions database is unreliable. The system either misses a genuine match — a compliance failure — or returns a volume of near-matches that must be cleared manually.

Accurate name resolution is therefore a precondition for screening, not a refinement of it. The mismatch takes recurring, script-specific forms, set out in Exhibit 2 on the following page.

EXHIBIT 2 · COMMON NAME-MATCHING FAILURES BY SCRIPT AND CONVENTION

SCRIPT / CONVENTION	WHY THE MATCH FAILS	EXAMPLE
Arabic & Kurdish	One name romanises several ways; the article "al" is joined, hyphenated or dropped; Kurdish characters do not map cleanly to Arabic	Mohammed / Muhammad / Mohammad as one person; Abd al-Rahman and Abdulrahman as one person
Spanish & Latin American	Two surnames, with systems retaining different parts and treating particles inconsistently	Juan Carlos Pérez García recorded as Juan Pérez, J. C. Pérez, or Pérez García
Greek	Inconsistent romanisation of Greek characters	Ιωάννης as Ioannis, Yiannis or Giannis
Chinese (Pinyin)	Distinct names collapse to one romanised string, and single syllables map to many characters	李 as Li or Lee; two unrelated people both as Wang Wei
Long Brazilian & Japanese	Multiple names and particles truncated or reordered, family-name-first ordering, and competing romanisations	Ana Beatriz de Oliveira Souza reordered; 大野 as Ono, Ohno or Ōno

Product impact. Poor name resolution produces two failure modes at once. A genuine match against a sanctions or PEP list can be missed because spellings do not align — a compliance failure. At the same time, loose matching generates false positives that enter a manual review queue, slowing onboarding and diverting analyst capacity from genuine risk to administrative clearance.

Why third-party-dependent platforms struggle. Using a third-party KYC provider built over an external matching engine, the logic that resolves names is not accessible to the product team. Support for a new script or naming convention depends on the upstream provider's development schedule.

How Shufti addresses it. Verifying identity documents across non-Latin scripts is the harder and more valuable capability, and it is a core strength of Shufti's owned stack. Once a document in Arabic, Kurdish, Greek, Chinese or Japanese is read and verified accurately, generating and reconciling the transliterated name is a comparatively basic layer — and because Shufti owns both the document verification and the matching logic, that layer is tuned in house. Adapting transliteration for a new script or convention is a scoped change the product team can order and receive on demand. For a platform built this way, these are not edge cases at all. They are the default.

Proof of Address verification and normalisation

The challenge for the product team. You are expanding into new markets where Proof of Address documents and address structures look nothing like the markets you already serve. Valid documents are being rejected, completion rates are slipping, and support tickets are rising. What you need is for the platform to read and normalise these formats now — not a wait while another vendor adds them.

Product impact. Document-upload requirements reduce completion rates, and rejections the customer cannot interpret generate support volume and abandonment. **Why third-party platforms struggle:** the address reading and normalisation logic sits inside the external engine, so support for a new format or document-free route depends on the provider. **How Shufti addresses it:** owning the reading and matching layer allows normalisation, special-character handling, and field standardisation to be adapted to the formats and markets served, and document-free Proof of Address to be routed through trusted data sources rather than manual review.

Why this matters. Proof of Address is a quiet source of onboarding attrition. Every document a customer must upload, and every rejection they cannot understand, costs you conversion. Expanding into a market with unfamiliar address documents should be a configuration task, not a development dependency that leads to bottlenecks.

Proof of Address extends the reconciliation problem to the address field. An address must be read from a document, standardised into a consistent form, and matched against the customer's stated address and the record on file. Further, the solution must be capable of accepting multiple Proof of Address documents — such as utility bills and bank statements — to ensure it is in the right domain. Each step can fail when the document is in another language, another format, or from a country the stack has not handled.

EXHIBIT 3 · PROOF OF ADDRESS FAILURE POINTS

FAILURE POINT	DESCRIPTION
Special characters and non-English scripts	A document in a local language rejected on characters the English equivalent would pass, indicating a limitation in the reading and matching layer rather than the document
Inconsistent formats across countries	Fields located and structured differently on each country's documents, requiring extraction and standardisation before comparison
Non-standard jurisdictions	Address proofs from countries the stack was not designed around, where plausibility must also be assessed — for example an address citing a floor that does not exist
Document-free verification	Confirmation of an address against trusted data sources rather than manual document review, requiring the address to be normalised and matched to a reliable source

Compliance geofencing and sanctions screening

The challenge for the product team. Compliance introduces a new geo-fencing rule. Specific issuing countries and nationalities must now be declined, the change has to be live before the next release, and a user routing through a VPN must not slip past it. It is a roadmap item the moment it arrives, and how quickly you ship it depends entirely on whether you can configure the rule yourself.

Why this matters

The rules for which users and documents you can accept move constantly, and they sit entirely outside your business. Get it too loose and you have a regulatory breach. Get it too tight and you block real customers and lose conversion. Either way, you are the one who has to keep the rules current, fast.

What changes

Compliance geofencing is the requirement to accept or reject a user based on location, origin, or document-issuing country, and to keep that decision correct as the rules change. The boundary is not fixed — it moves with sanctions regimes, licensing, and local law. The rule set is layered: blocking regions where a product is unlicensed or illegal, applying different legal ages by region, preventing a user who registers in a permitted location from operating from a prohibited one, and refusing documents issued by countries on a restricted list, with OFAC, EU, UN, and local lists revised regularly. Many regulated businesses also screen against a local or central-bank blacklist alongside the global lists.

Why a single feature is not enough for geo-location compliance in AML

Location signals are insufficient on their own. A user in a sanctioned jurisdiction can route a connection through a permitted country using a VPN or proxy. If the control stops at IP geolocation, that user passes. The reliable control operates at the identity layer: reading the document's issuing country, identifying that it falls under a restricted list, and declining that document regardless of the apparent connection origin. A sound rule combines connection location, issuing country, the nationality shown on the document, and VPN detection. The same geolocation feature is then used in implementing age verification.

A sequence you will recognise

A sanctions list is revised late on a Friday. Compliance requires the change live by Monday. On a third-party-dependent platform, the change becomes a support request in a queue you do not control, an engineer is reassigned from the sprint to build an interim measure, the release slips, and by mid-week revenue is asking why conversion fell in a market you over-blocked to remain compliant. The underlying change was a policy edit. The cost was a sprint.

EXHIBIT 4 · IMPLEMENTING A SANCTIONS OR GEO CHANGE: TWO CYCLES

List or rule changes	List or rule changes
Raise a support request	Rules maintained and updated centrally
Wait in the provider backlog	Configure market and risk rules directly
Assign engineering to a workaround	No engineering cycle required
Release delayed, conversion affected	Change effective, roadmap unaffected
Why third-party-dependent platforms struggle. The geographic and sanctions logic is set upstream, so each list revision or rule change is a dependency on the provider rather than a configuration the product team controls.	

How Shufti addresses it

Because Shufti owns its verification stack, these changes can be implemented quickly without waiting for an upstream provider — whether skipping SSN verification for non-US documents, declining documents from sanctioned jurisdictions, introducing IP and VPN checks, or adapting market-specific rules. For product teams, compliance changes become product configurations rather than engineering projects.



MacroPay's Multi-Market Compliance, Powered by Configurable Verification

Macropay, a fintech building payment infrastructure across the EEA and additional international markets, partnered with Shufti to establish a robust and adaptable compliance framework without increasing operational overhead. It implemented Shufti as a centralised identity verification and compliance layer, unifying KYC and KYB across multiple regions and regulatory environments. One pressure Macropay faced was "tight timelines driven by both regulatory requirements and business growth" — the need to deliver compliance changes on schedule.

“

Shufti gives us the flexibility and coverage we need to build a compliant, multi-market payment setup without overcomplicating the product.

Flavia Mic
Product Manager

”

Data residency and in-country deployment

The trigger for your team.

You are ready to enter a new market when compliance flags that its regulator requires verification data to remain in-country, or the verification software to run inside your own environment. This is not a setting you can switch on. If your provider is cloud-only, the market stays shut.

Why this matters

Data residency is the requirement most likely to close a market outright, because it cannot be solved in the product layer. If your platform is cloud-only, a sovereignty rule is not a feature gap — it is a locked door, and the market does not open.

What changes

Data residency becomes a product constraint when regulations require that verification data remain within a specific country, or that verification software run within the customer's own environment rather than the vendor's cloud. This is a deployment requirement that a standard cloud service cannot satisfy through configuration alone.

These requirements are typically most stringent in highly regulated sectors such as banking. Financial institutions operating under national central bank or similar frameworks may be prohibited from using cloud-based verification services. Instead, verification must run on-premise within the institution's own infrastructure, verification data must remain within national borders, and outbound internet access from the on-premise system may be restricted. Similar in-country storage requirements apply to other regulated entities such as payments providers. More broadly, the same pattern arises wherever data-sovereignty regulations apply, including environments that require fully offline operation or deployment within customer-managed, in-country data centres.

Why third-party platforms struggle

A provider without control of its own deployment model cannot relocate where the software runs or where data is stored, so a sovereignty rule falls outside what it can offer.

How Shufti addresses it

meeting data residency and sovereignty rules comes from flexible deployment options. Shufti offers in-country cloud, on-premise, offline, and SaaS deployment, so the model can be matched to whatever a jurisdiction requires rather than being limited to a single shared cloud. A data-residency requirement becomes a deployment option rather than a closed market.

Owning the tech stack: how Shufti has helped compliance and product teams

Real-world identity verification rarely fails because of a single universal rule. It requires continuous customisation as governments modernise administrative systems, countries follow different cultural naming conventions, and local scripts demand specialised OCR models. The examples below illustrate the localisation and configuration that become essential when operating across diverse jurisdictions.

VIETNAM

Address modernisation

Old ID cards carry outdated address conventions after Vietnam's July 2025 administrative reforms. Shufti recognises the card and returns the updated convention automatically, so a genuine user is not rejected over formatting.

JAPAN

Name ordering by spacing

Japanese cards write the family name first, breaking name matching. Shufti splits the name based on the spacing on the card, so first and last names resolve correctly.

CHINA

Generational surname

The generational family name sits where most systems expect the first name. Shufti extracts it into the correct field every time, removing a systematic source of false rejections.

BURMA

In-house OCR

Generic OCR reads Burmese poorly and rejects genuine users. Shufti's own Burmese model reaches roughly 85% accuracy — 20 to 30% above a general-purpose engine.¹

The same pattern extends on demand to Hong Kong, Hungary, the Gulf states, Arabic-script documents, and native-to-Latin transliteration — each a scoped change rather than an edge case deferred to a vendor. Across its verification, onboarding, back-office, and analyst layers, Shufti reports more than 1,800 tailored configurations shipped.

How A Global Social Media Platform Dismantled Organised Fraud Rings Through Owned-Stack Fraud Detection

New fraud patterns are one of the requirements that lands on a Product Owner without warning. On the same Japanese social media platform, Shufti detected industrial-scale fraud rings, evidenced by multiple IDs submitted against identical backgrounds, and responded with Repeat Offender Suppression built on digital fingerprinting and behavioural tracking, alongside replay and stream-tampering detection to counter video manipulation. Because the stack is owned, these countermeasures were engineered as the attack evolved rather than waited on.

¹ Accuracy figure based on Shufti internal benchmarking against a general-purpose OCR engine. Source: Shufti global social media platform case study.

How Shufti's Owned Tech Stack has Empowered Product Teams

New-market expansion and evolving compliance requirements demand rapid KYC adaptability. Shufti delivers market-specific customisation at speed. This helps digital-first businesses prevent operational bottlenecks.

CAPABILITY	CLIENT CHALLENGE → HOW OWNING THE STACK RESOLVES IT
Vietnam name splitting	The full name does not map to first and last name fields → Shufti splits and, on specific card types, correctly orders the name fields
Vietnam permanent cards	Cards with no expiry date fail as expired or empty → Shufti recognises permanent designs and sets a valid non-expiring date so they pass
Hong Kong & Hungary names	Each country structures names differently, by commas, hyphens, or order → Country-specific splitting rules are shipped on demand
Arabic-script capture	Names and data must be retained in native Arabic, not forced to English → Shufti preserves native Arabic and formats per country as required
Native-to-Latin transliteration	Document numbers in native script will not reconcile with Latin records → Shufti transliterates the document number to Latin
Gulf full-name handling	Gulf-state IDs split names unreliably → Shufti prioritises the full-name field for Gulf documents
Region-tuned face matching	One global face-match threshold over-rejects in some regions → Shufti sets similarity thresholds per region to balance pass rate and security
Risk-based age step-up	Running full document checks on everyone adds friction → Shufti uses risk based age estimation
AML filter configuration	Screening must cover the right lists, across sanctions, PEP, and adverse media → Shufti applies configurable filters as per client risk exposure
Sanctions-list transparency	Compliance needs assurance the sanctions data is current → Shufti shows clients when the sanctions list, such as OFAC, was last updated
Blacklist review workflow	Auto-declining on a blacklist hit risks wrong rejections → A hit can be routed to analyst review before any automated decision
Geo and country control	Issuing-country rules must apply even to trial accounts → Shufti accepts or rejects by location, issuing country, and nationality, trials included
Capture-mode control	Some flows need camera-only, others must allow upload → Shufti toggles online and offline capture per client and per service
Video KYC flexibility	Regulated flows and bespoke VKYC scripts vary by client → Shufti runs jurisdiction-specific VKYC journeys and custom on-screen instructions
Proof of Address flexibility	Rigid Proof of Address rules reject valid documents and drive abandonment → Shufti accepts a wider range of proof of address documents
Deployment and residency	Regulators require data in-country or software on-premise → Shufti deploys on-premise, offline, or in-country, to help comply with data residency requirements



VERIFY ACROSS
BORDERS.

An evaluation framework for product leaders

Use a proper evaluation framework to pressure-test any verification platform before you commit. The single question that predicts your future roadmap load is whether the provider owns its stack, because ownership decides whether change is configuration or a support ticket — and how quickly matching, geofencing and deployment can take place.

Vendor evaluation guidelines

The objective of vendor evaluation should extend beyond assessing current feature availability. It should determine whether the vendor's platform can continuously adapt to changing regulatory requirements, support expansion into new markets, and enable the business to meet its compliance obligations without repeated platform replacement or extensive redevelopment.

COMPLIANCE ADAPTABILITY AND TECHNOLOGY OWNERSHIP

- Does the vendor clearly demonstrate the value of its owned technology stack in responding to evolving regulatory and compliance requirements?
- Does the vendor offer rapid customization to meet the business's specific operational and compliance requirements without lengthy development cycles?
- Does the vendor provide these customizations through its own technology and engineering capabilities rather than outsourcing them to third-party providers?
- Can new compliance requirements be implemented through configuration or controlled product enhancements rather than bespoke development?
- Is the provider transparent about which capabilities are configurable today and which remain on the product roadmap?
- Is a genuinely new regulatory requirement delivered through a defined implementation process with predictable timelines rather than an open-ended product request?

MARKET EXPANSION AND REGULATORY COVERAGE

- Is the platform sufficiently configurable to support expansion into new countries and regulatory jurisdictions without requiring a new solution?
- Can verification policies be adapted to meet country-specific regulatory requirements while maintaining a consistent operating model?
- Can name matching and transliteration be tuned to support new languages, scripts, and local naming conventions, without vendor engineering involvement?
- Can address extraction, parsing, and normalisation be adapted to new address formats and markets, including document-free Proof of Address verification?

IDENTITY VERIFICATION COVERAGE

- Does the solution extend beyond document verification to support document-free identity verification through government eIDs, authoritative databases, and trusted commercial data sources?
- Does the provider demonstrate an ongoing strategy to expand its coverage of government eID and authoritative database integrations across new markets?

RISK MANAGEMENT AND DECISION FLEXIBILITY

- Will expanding verification coverage help increase customer acquisition by enabling users who cannot easily provide identity documents?
- Can verification workflows, decision rules, and acceptance thresholds be configured to align with the organisation's risk appetite and exposure?
- Is the vendor sufficiently flexible to support a risk-based approach, allowing verification requirements, controls, and CDD measures to vary according to assessed risk?
- Can risk policies be adjusted independently for different products, customer segments, jurisdictions, or regulatory obligations?
- Can geo-restrictions, sanctions screening rules, and country risk policies be reconfigured as regulations change, and are sanctions and watchlists actively maintained by the provider?
- Does compliance geofencing combine multiple signals — user location, document issuing country, nationality, and other attributes — to prevent prohibited users bypassing controls through IP masking or VPNs?

DATA SOVEREIGNTY AND DEPLOYMENT FLEXIBILITY

- Does the platform support flexible deployment models — cloud, on-premise, fully offline, and customer-hosted in-country — where data sovereignty regulations require them?
- Does the provider support compliance with data residency and sovereignty requirements through these deployment options?
- Can verification data remain entirely within a required jurisdiction without compromising platform functionality?

LONG-TERM PRODUCT VIABILITY

- Does the vendor demonstrate continuous investment in expanding compliance capabilities, verification methods, and regulatory coverage?
- Is the platform designed to evolve alongside changing regulations rather than supporting only current compliance requirements?
- Will the solution remain suitable as the organisation expands into new markets, adopts new products, or faces new regulatory obligations?

Conclusion

Regulatory change in KYC and AML is continuous, and the burden of implementing it falls on product teams. The architecture of the platform determines whether that burden is manageable. A platform assembled from third-party engines converts each regulatory change into an engineering dependency and a wait. A provider that owns its identity verification stack converts the same change into a configuration task or a scoped internal build.

Compliance requirements evolve continuously, with greater complexity as businesses enter new markets. Shufti enables identity and KYC product owners to configure risk rules around local regulatory and operational requirements without creating development bottlenecks.

This includes processing complex scripts, reconciling names across databases, adapting to administrative reforms, and setting thresholds aligned with sanctions and AML obligations. The result is faster market expansion with greater control over compliance workflows.

Shufti builds and owns its complete identity verification stack in house — name matching, biometrics, OCR, and screening — without dependency on a third-party engine. The cases that constrain a resold platform (an unfamiliar naming convention, a revised country restriction, a newly required document, or a data-residency mandate) are standard operating conditions for an owned stack. **Their edge cases are our default.**

[Talk to Shufti](#)