Regula

Identity Verification in the Age of AI Agents

How Organizations Adapt Identity Systems to Fraud, Automation, and Machine-Driven Interactions

Regula Global Research Report | 2026

Table of Contents

Regula

“

Identity verification was built for a world where every interaction involved a real person. That assumption no longer holds. Today, organizations must not only verify identity signals — they must understand whether the interaction itself is genuine, live, and human-controlled.

Documents can be manipulated at scale. Faces can be generated or injected into camera feeds. Automated actors can move through onboarding, login, and transaction flows designed for humans. The result is a widening gap between how identity systems were built to operate and the environment in which they now function.

This report examines how organizations are adapting: where controls exist, where confidence is limited, and what it takes to maintain trust when identity signals can no longer be trusted by default.

Henry Patishman,
Executive VP of Identity Verification Solutions,
Regula

Methodology

- 850 decision-makers in fraud detection and financial crime

- 7 markets:

UK

US

UAE

Germany

Mexico

Brazil

Singapore

- Industries include:
  - Banking
  - Government
  - Financial Services
  - Crypto
  - Telecommunications
  - Gaming/Gambling

- Results are accurate to ±3.4% at 95% confidence

- Fieldwork conducted by Sapio Research, March 2026

From Threat Signals to What's Next

Part I of this report established the core threat signals. Part II examines the business impact, control gaps, and future readiness.

What Part I Established

98%
concern about identity-related threats.

87%
AI-assisted or automated actors attempting to pass identity processes.

69%
AI tools are present in identity flows.

35%
AI-generated impersonation is a major concern.

26%
machine-operated actors acting on behalf of users.

What Part II Examines

Impact of failures.
The financial, operational, and reputational cost of identity failure.

Readiness gaps.
Where controls, visibility, and policy still fall short today.

Decision accountability.
How to justify, explain, and improve identity decisions at scale.

Key Findings

Confidence is limited	48%	of organizations consider technical controls reliable.
Liveness remains a blind spot	52%	cannot fully verify that biometri data was captured live.
Synthetic fraud's low visibility	41%	cannot fully assess whether identity signals are manipulate
Decision explainability lags	50%	can trace identity decisions end to end.
Regulatory scrutiny is here	82%	have been required to justify identity decisions externally.
Strategy is catching up	47%	explicitly address AI-assisted interactions in their identity strategy.

Trust Signals

Built for Humans. Stress-Tested by Machines.

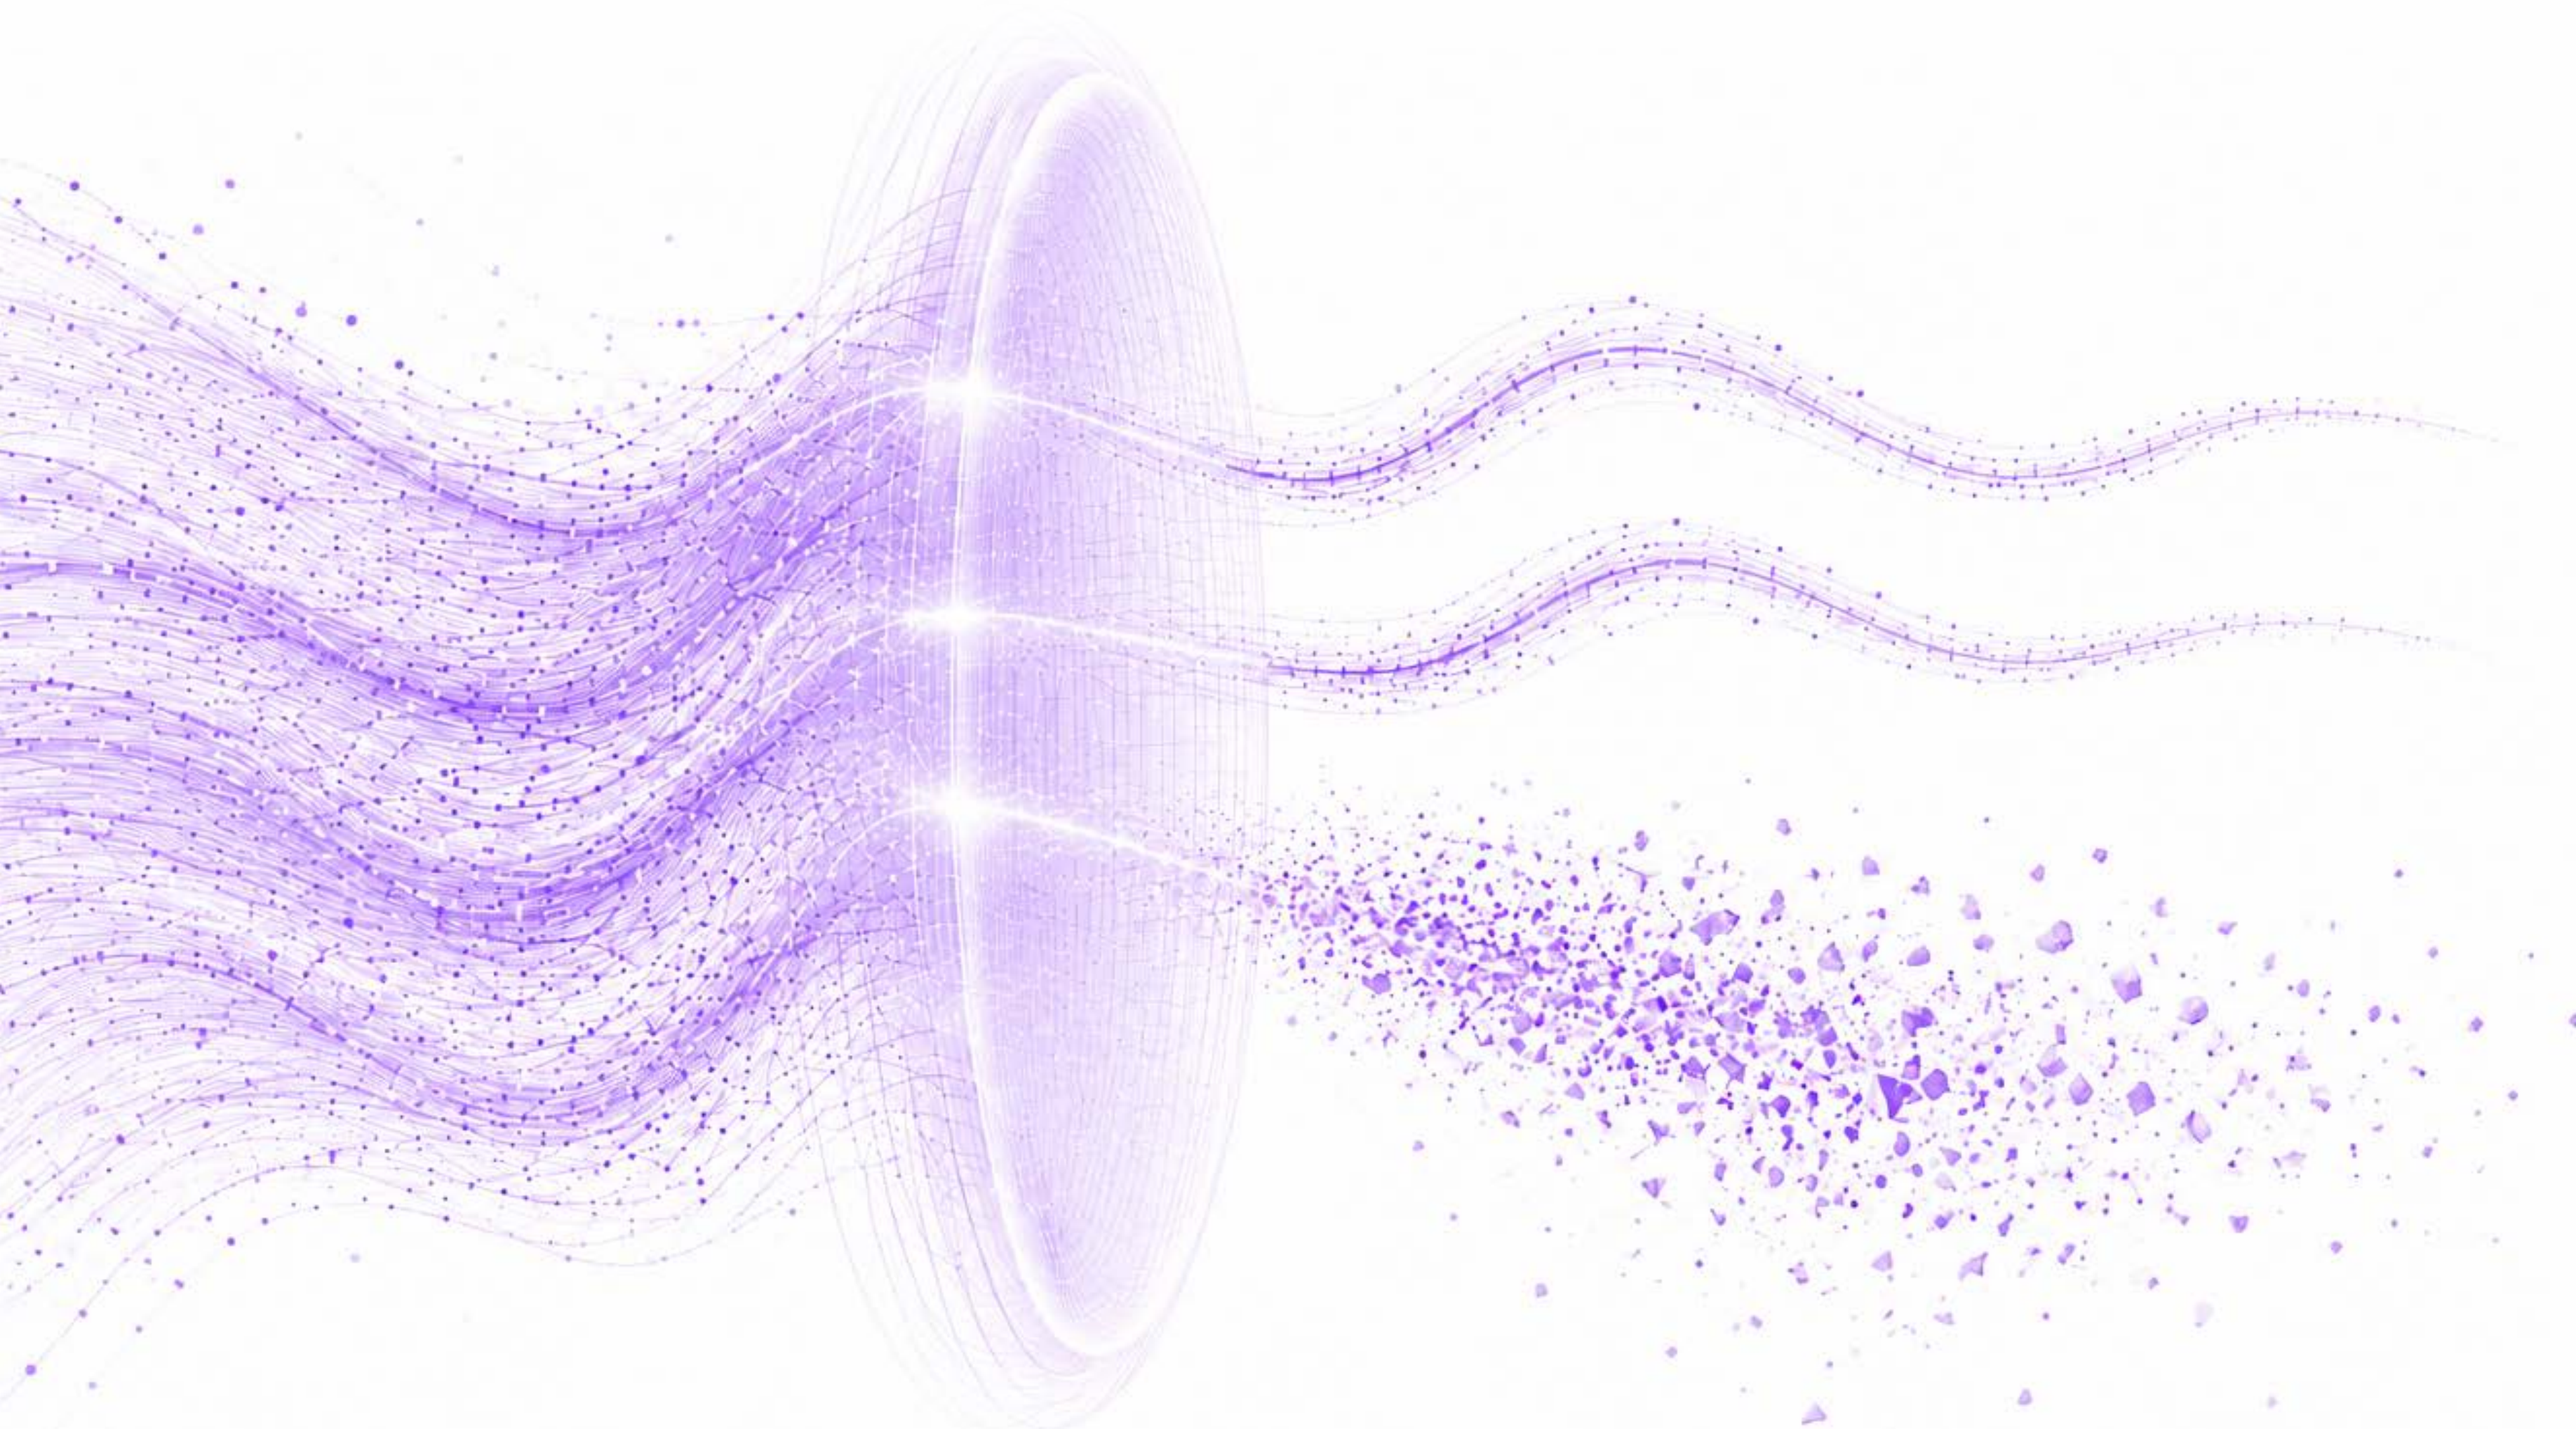

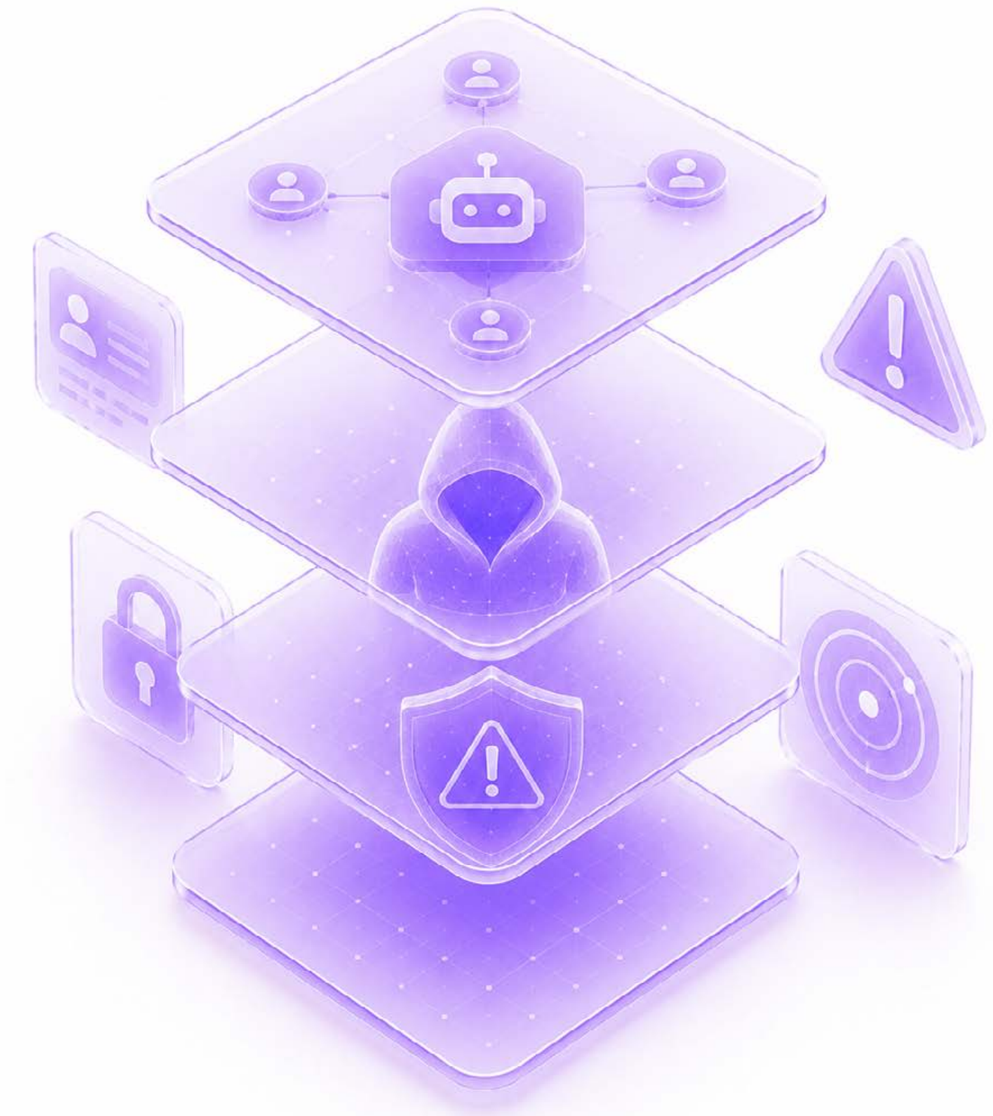

AI Agents in Identity Verification Are Not a Future Scenario

AI agents become a practical identity verification challenge. In identity workflows, AI agents are not only chatbots or productivity tools. They can appear as automated or AI-assisted actors moving through onboarding, login, recovery, support, or transaction flows that were designed for humans. They may use generated documents, synthetic faces, manipulated camera feeds, scripted behavior, or legitimate identity fragments to pass checks.

The survey data shows many organizations have identity checks in place, but the harder question is whether those checks can still prove trust in an AI-driven environment: Is the document real? Was the biometric captured live? Do the documents and biometrics belong to the person behind this session? Is there even a real person behind the session altogether?

Document Verification Methods

Data cross-validation is the most trusted document verification method.

By comparing identity information across multiple independent sources, it helps detect inconsistencies that a single document check may miss.

This reflects a broader shift in identity verification. A document, face, or database match may look correct on its own, but AI-driven fraud can combine valid-looking signals in misleading ways. Trust increasingly depends on how document, biometric, database, device, session, and risk signals support each other across the identity flow.

↦ **Key takeaway:**
Identity trust depends on connected evidence, not any one signal alone.

Most trusted document verification methods

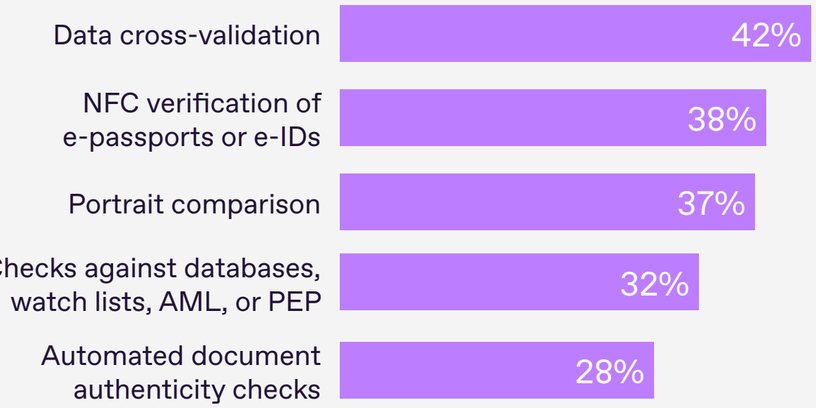

Trust is concentrated around a narrow set of verification signals.

Synthetic Content Detection

41%

cannot confidently assess whether identity evidence is authentic or synthetic.

AI-generated images, manipulated documents, synthetic biometric data, and altered media may not always look suspicious. If these inputs are not detected, they can enter the identity workflow as accepted evidence and influence real decisions.

This matters for AI agents because they do not always need to “break” an identity system. They may only need to submit convincing evidence through a normal user flow and avoid detection.

↦ **Key takeaway:**
Fake identity evidence can become part of real decisions.

Synthetic content detection capabilities

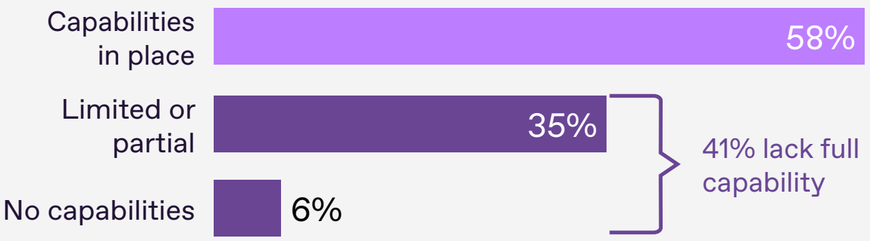

Detection is present in many organizations, but four in ten still lack full authenticity assessment.

Proving Human Presence

76%

have controls to verify human presence — but only 48% consider them reliable.

Organizations need to prove that a real, live person is behind the interaction — not a bot, deepfake, replayed video, injected camera feed, or automated process.

This is the clearest way to explain the AI agent threat: systems built to verify people increasingly encounter activity that behaves like a user, submits identity evidence like a user, and moves through digital journeys like a user. But that does not always mean a real human is present.

↦ **Key takeaway:**
Human presence is becoming one of the most important trust signals in remote identity verification.

Human presence verification capability

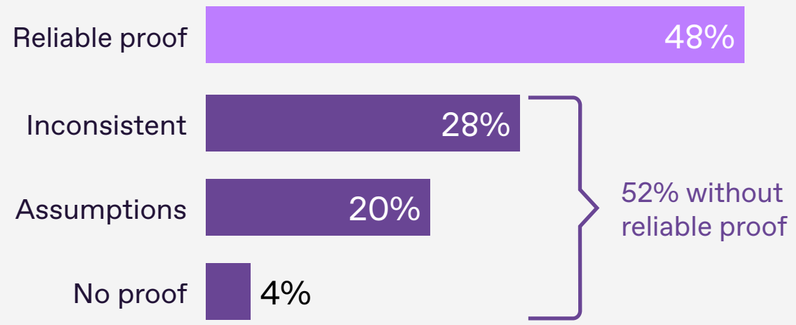

Only 48% have reliable controls; nearly one in four rely on assumptions or cannot prove human presence.

Human Presence Assurance Gap

Confidence is uneven across markets and industries.

Organizations that depend on remote identity interactions do not all have the same ability to prove that a real person is present.

Reliable human presence controls

By country

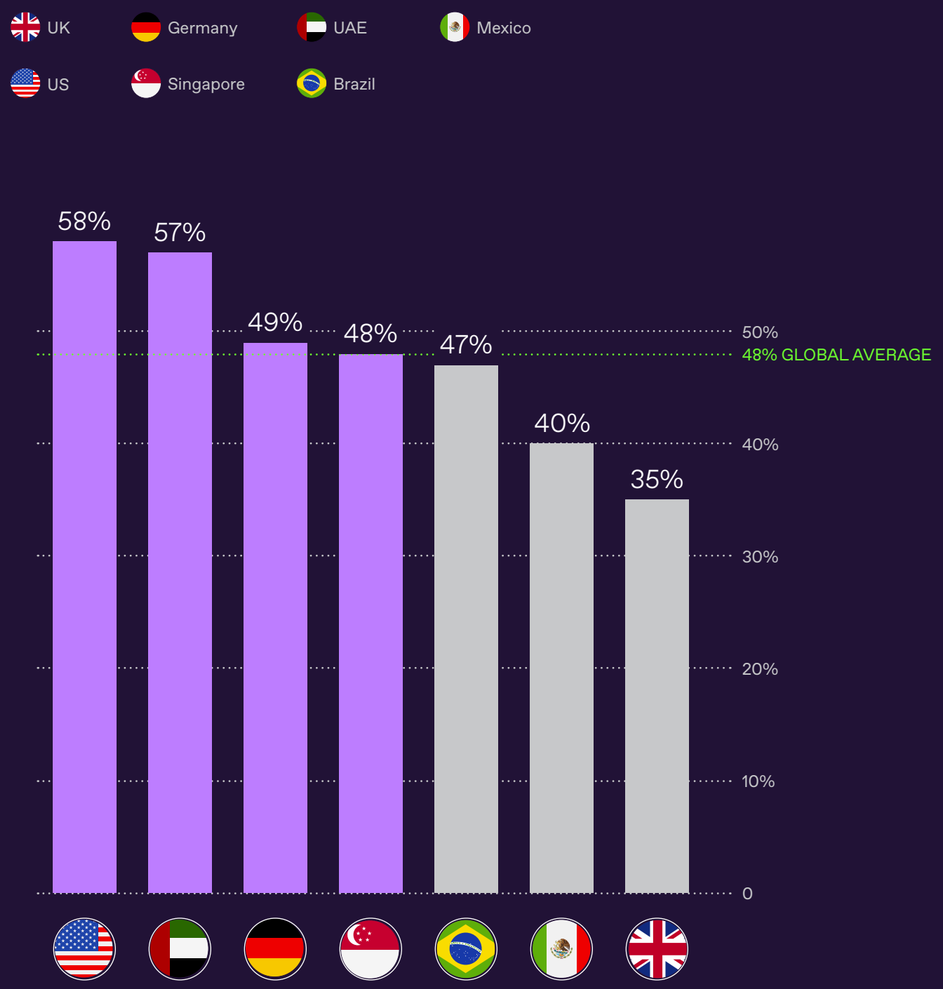

Confidence in human presence controls is uneven across markets, suggesting different levels of technical assurance.

By industry

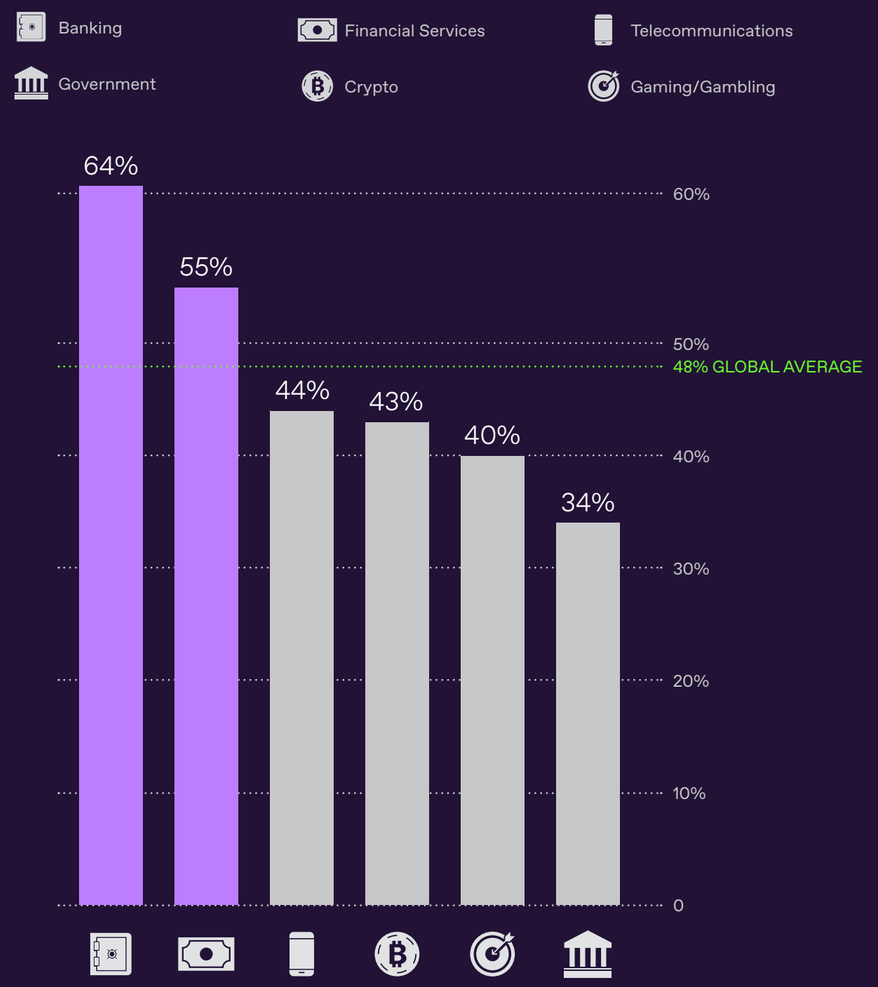

Human presence assurance is strongest in regulated financial sectors, while industries with fragmented or high-volume identity journeys show lower confidence.

Liveness Verification Remains a Blind Spot

52%

cannot fully verify that biometric data was captured live.

Liveness verification helps confirm that biometric data — such as a selfie or face scan — is captured from a real person in real time.

This is increasingly important because biometric evidence can be manipulated. Attackers may use prerecorded videos, synthetic faces, deepfakes, or injected camera streams that appear legitimate to the system.

A face match can show similarity. Liveness helps prove reality.

↦ **Key takeaway:**
Face matching is not enough when biometric data can be replayed or injected.

Biometric liveness detection capabilities

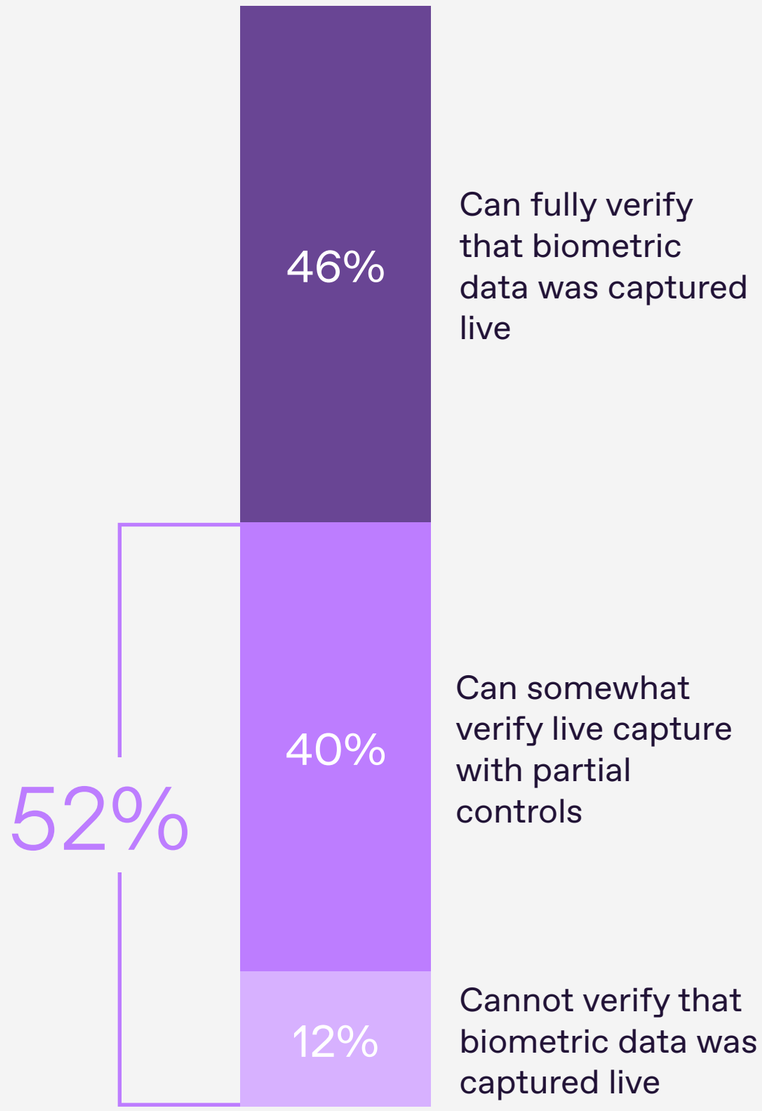

More than half of organizations lack full biometric liveness assurance.

Liveness Assurance Gap

Full live biometric capture is not evenly proven.

Liveness is becoming a core defense against AI-assisted identity activity, but the ability to fully verify biometric live capture varies sharply across markets and sectors.

Fully verify biometric live capture

By country

UK · Germany · UAE · Mexico · US · Singapore · Brazil

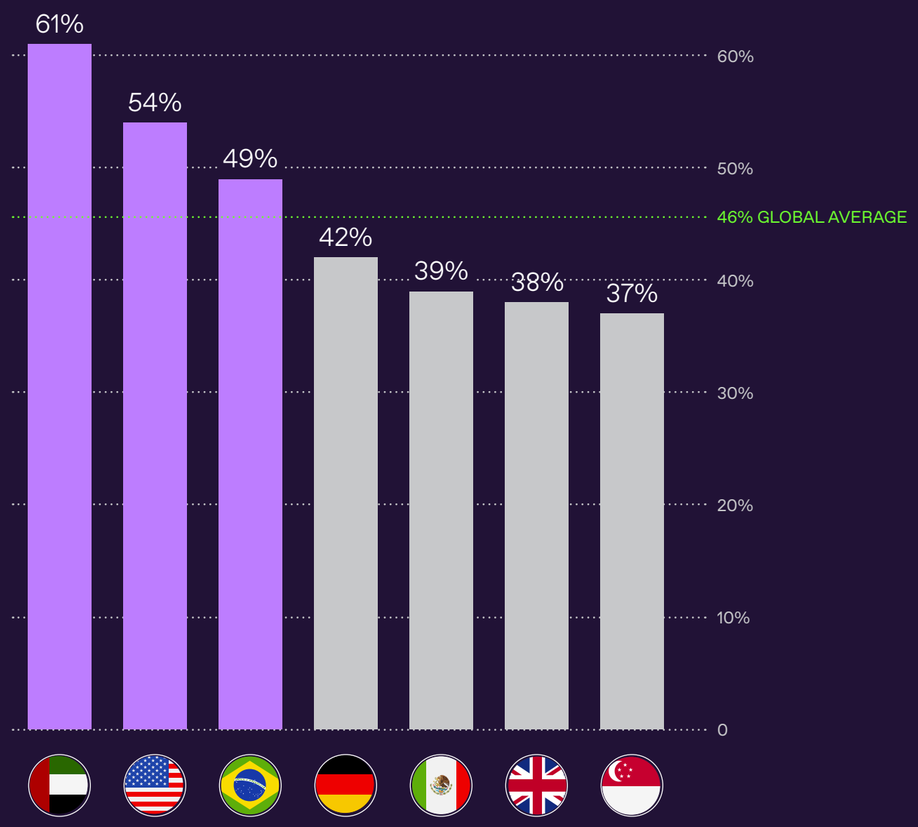

Full liveness assurance varies by geography, showing uneven confidence in proving live capture.

By industry

Banking · Financial Services · Telecommunications · Government · Crypto · Gaming/Gambling

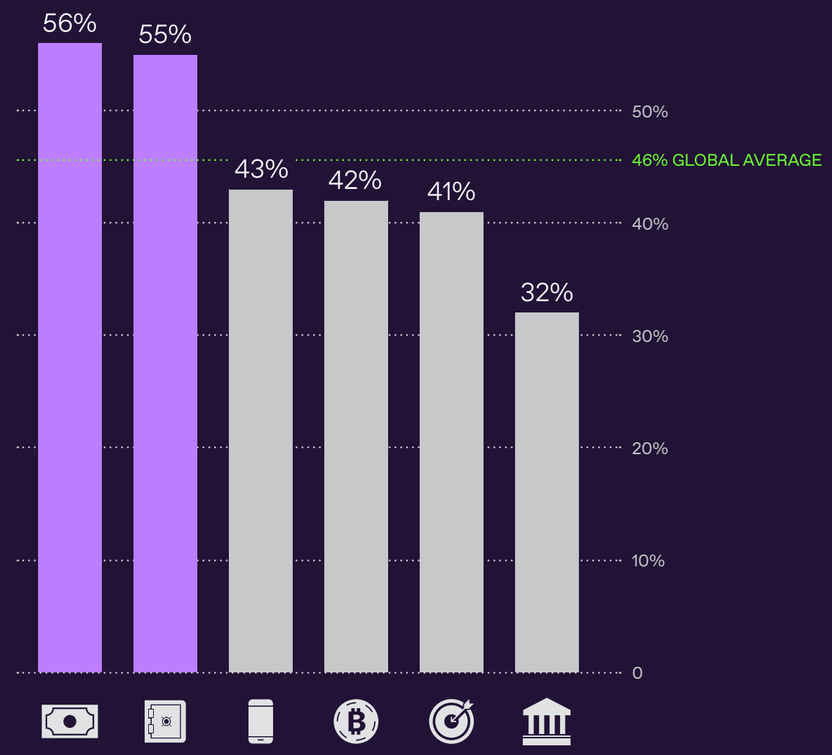

Regulated financial sectors lead; Government reports the lowest confidence.

Operationally Fragmented Systems

AI-driven identity risk is harder to detect when signals are scattered.

Most organizations do not operate a single, unified identity verification environment. Workflows are spread across vendors, internal systems, regional deployments, and specialized tools — making it harder to connect document, biometric, fraud, session, and decision signals. AI-driven activity often becomes visible only when signals are correlated across systems.

↦ **Key takeaway:**
Fragmented IDV environments reduce visibility and make coordinated identity attacks harder to detect.

Current IDV operating models by operating model type

Unified

21% Single vendor

11% Platform approach

Hybrid

25% Mix of vendors and in-house components

21% Multiple IDV solutions for regions, countries, services

13% Additional vendor for measurement

Fragmented

8% Highly fragmented

Business Impact and Future Readiness

Identity failures already affect the business.
Response is still maturing.

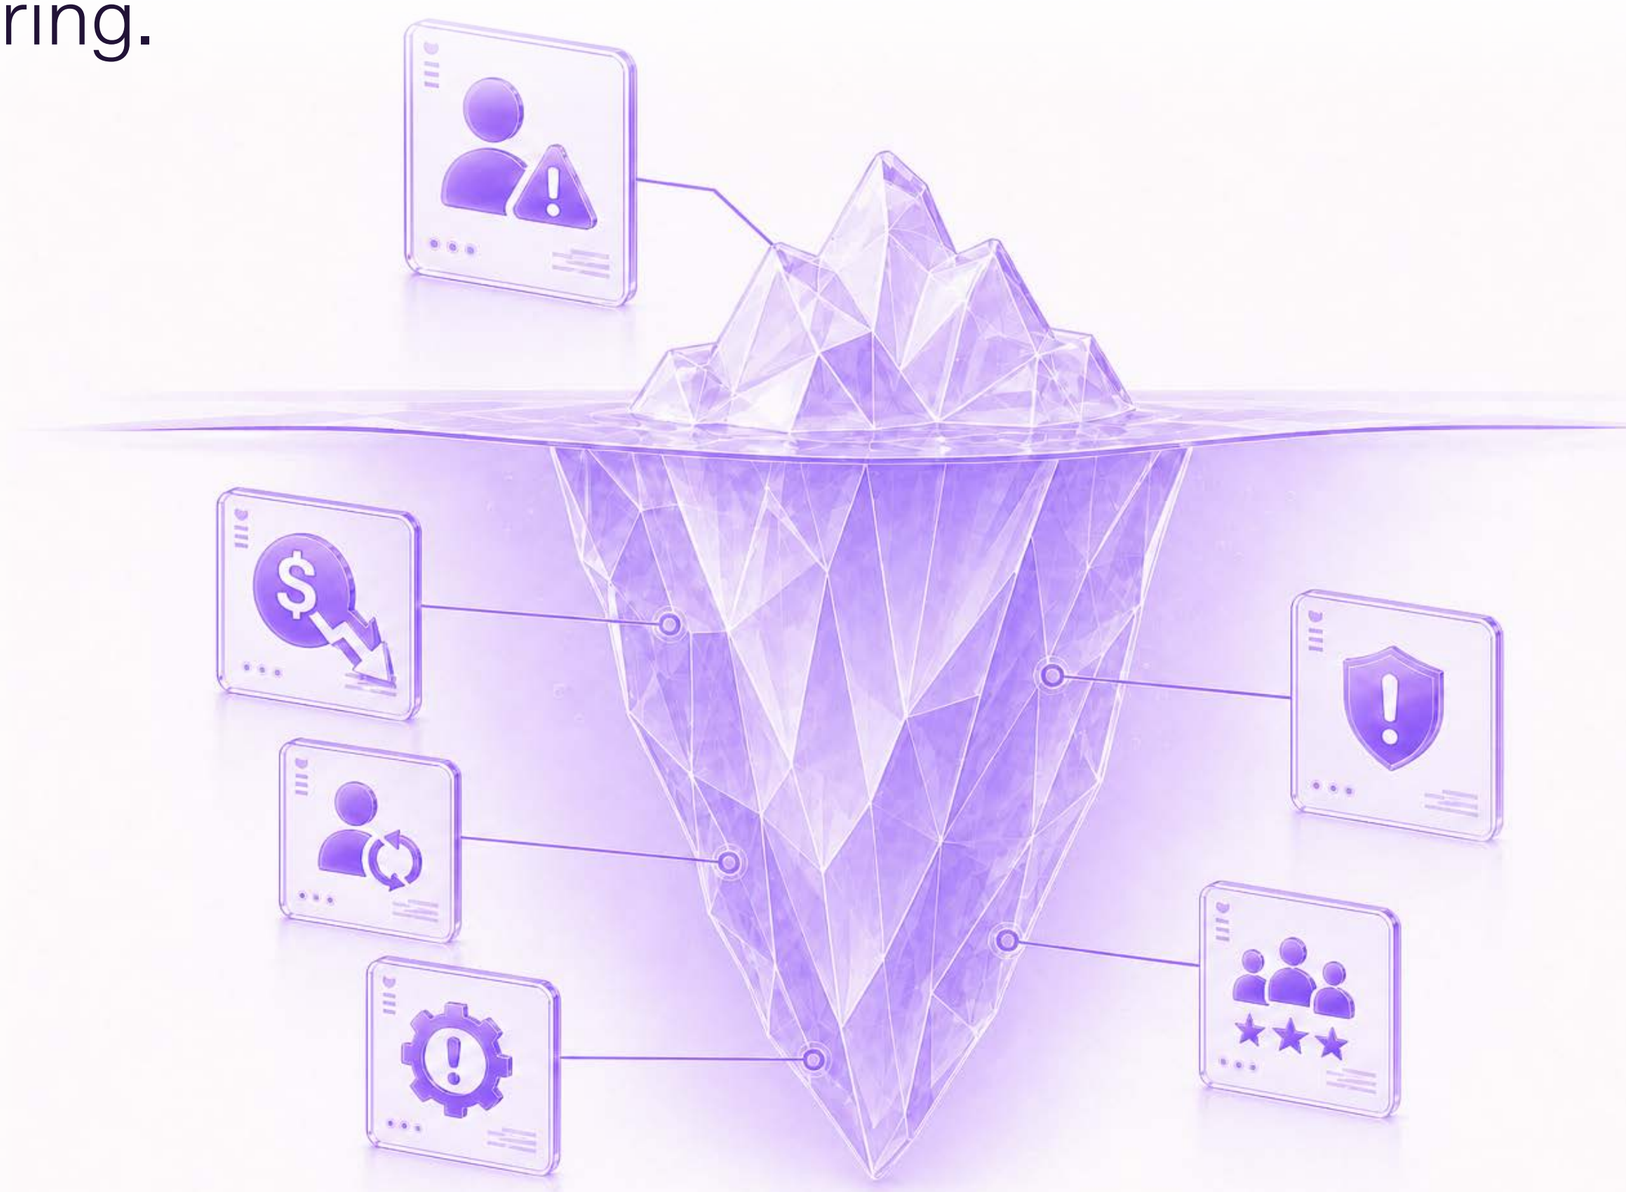

Enterprise-Wide Impact

92%

report business impact from incorrect identity verification results.

When identity decisions are wrong, the consequences can affect revenue, compliance, operations, customer trust, and brand reputation at the same time.

AI-assisted interactions increase the pressure. A failed identity decision may reveal that an organization cannot reliably determine whether an interaction was real, live, human-controlled, or manipulated.

The readiness challenge is clear: organizations recognize the threat, but response remains uneven.

↦ **Key takeaway:**
Identity verification is becoming a business resilience issue, not only a security control.

Impact of incorrect identity verification results

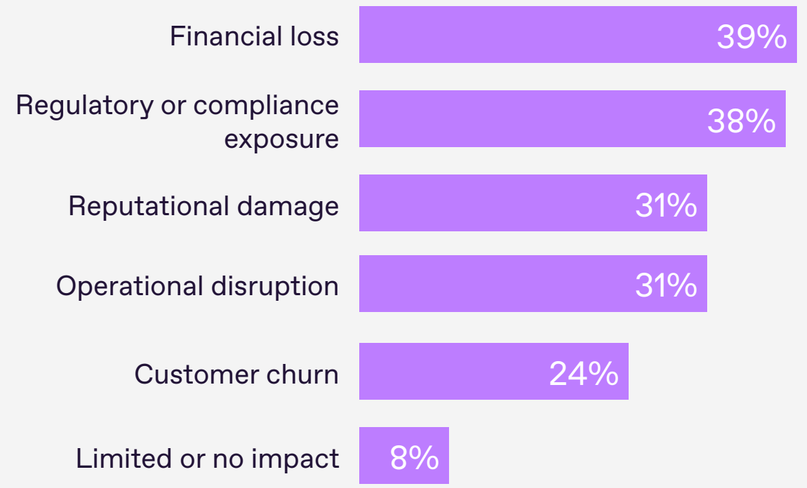

Identity failures create consequences across finance, compliance, operations, customer trust, and reputation.

AI Suspected, Core Evidence Holds

AI suspicion does not change what organizations trust most.

Organizations rely on almost the same identity tools and signals in routine digital verification and in post-incident investigations involving suspected AI-manipulated or synthetic identity interactions.

Biometrics remain the top trusted signal in both contexts, while chip/RFID-enabled government documents stay second. The shift is minimal: 0 pp for biometrics, and +1pp for chip/RFID documents.

↦ **Key takeaway:**
Suspected AI manipulation does not move organizations away from core identity evidence. It raises the bar for how reliably these signals are captured, verified, and combined.

Pre- vs. post-incident trust

Despite the AI-related investigation context, the hierarchy of trusted evidence remains largely unchanged.

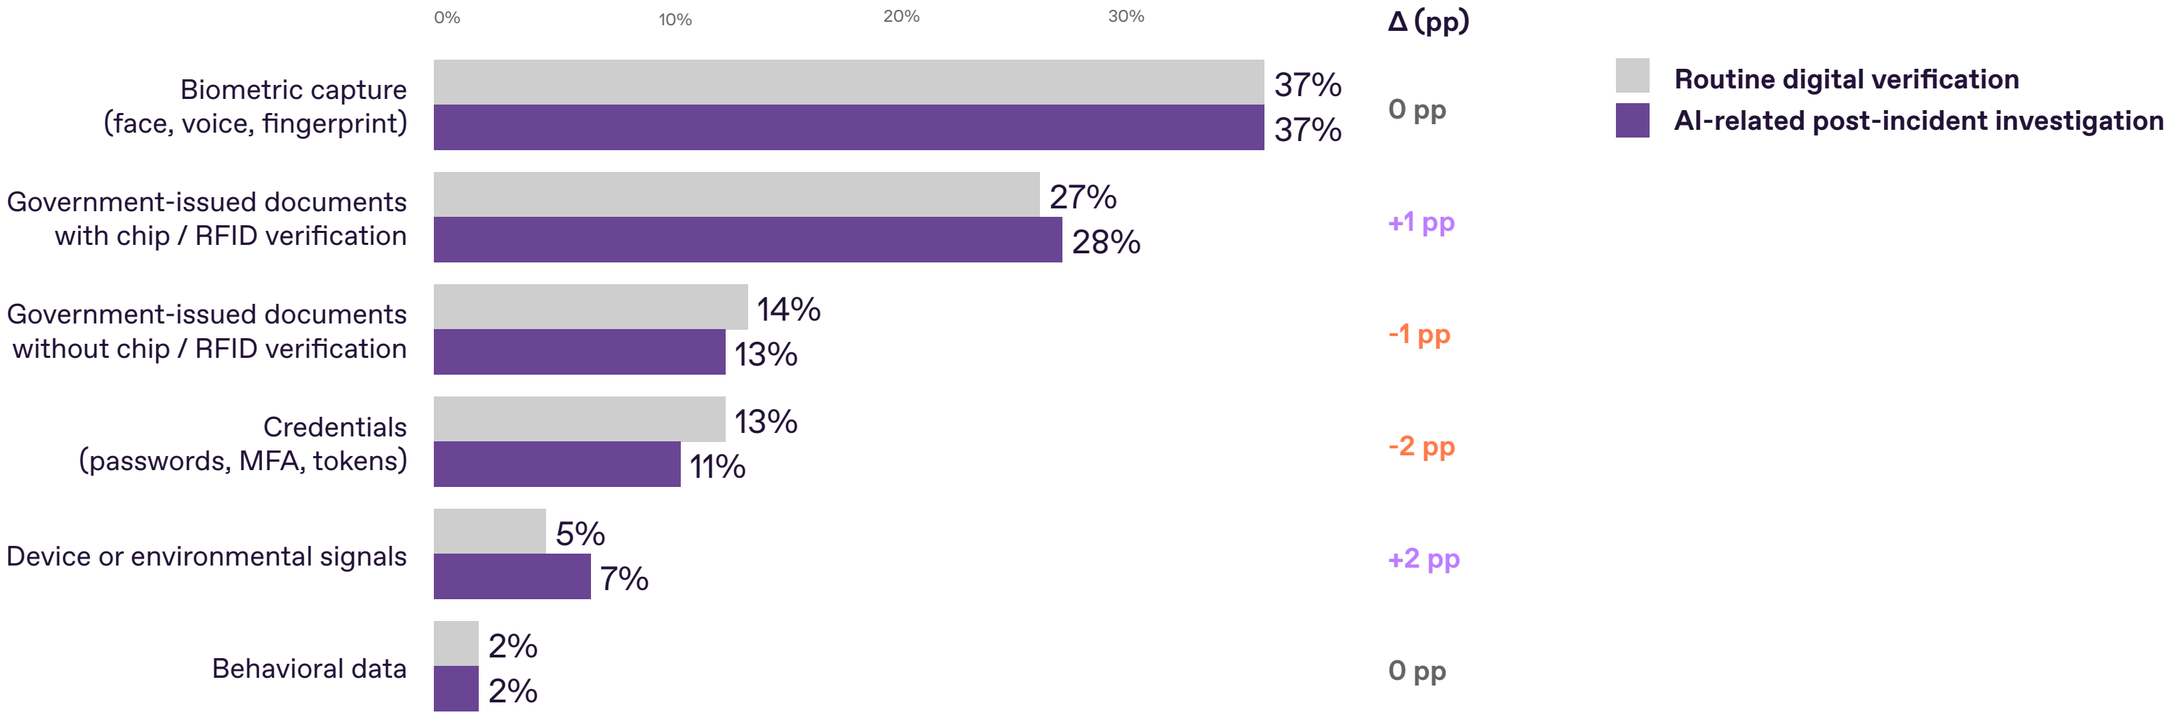

Strategy Is Being Updated

47%

explicitly address AI-assisted interactions in their identity strategy.

Many organizations now recognize AI-assisted identity activity as an operational challenge. They are updating strategies with stronger fraud controls, liveness detection, synthetic content detection, behavioral monitoring, and expanded governance.

But many remain in partial or planning stages, still defining ownership, controls, escalation rules, and evidence requirements.

↦ **Key takeaway:**
Organizations increasingly recognize AI identity risk, but many are still building the controls and governance to respond consistently.

Strategy for AI-assisted interactions

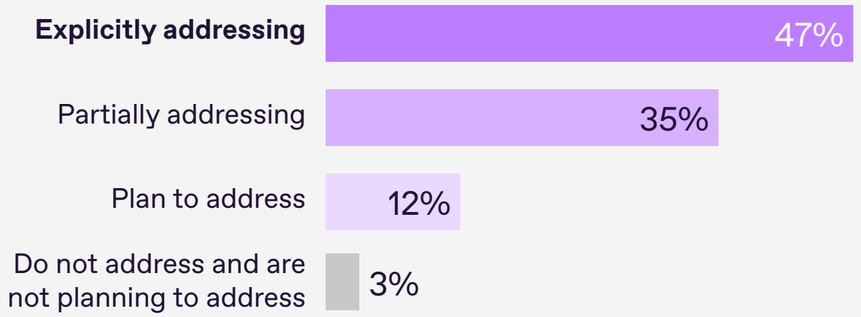

AI-driven identity risk is entering strategy, but many organizations remain in partial or planning mode.

AI Identity Strategy Is Advancing Unevenly

Some markets and sectors are moving faster from awareness to action.

Many organizations recognize AI-assisted identity activity as an operational challenge — but recognition does not always translate into strategy.

Explicit AI identity strategy adoption

By country

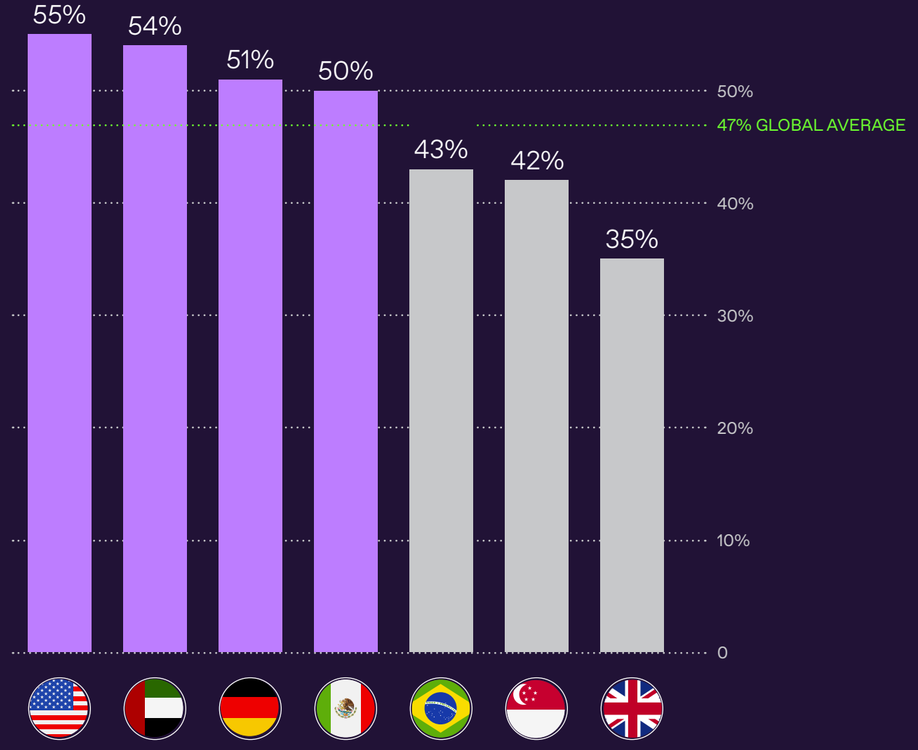

Adoption varies across markets, suggesting different levels of readiness.

By industry

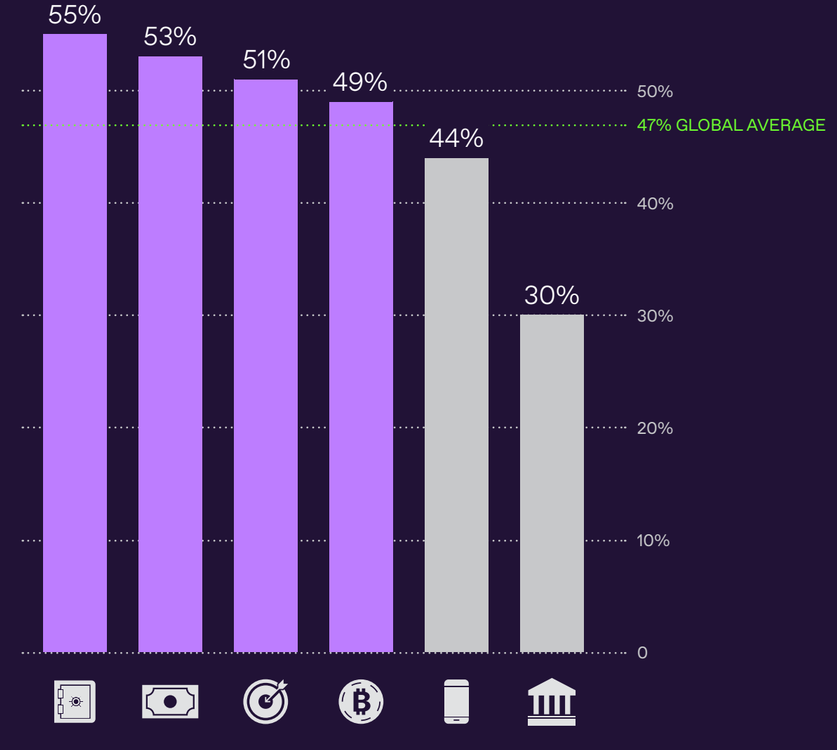

Banking and Financial Services lead; Government is lowest.

Future Priorities

Organizations are prioritizing live interaction controls.

Organizations are focused on capabilities that help detect AI-driven identity threats during live interactions. Top priorities include verifying biometric capture and human presence, triggering additional checks based on risk, and correlating document, biometric, and session data.

These capabilities help answer the most urgent question: is this a real person, acting live, with evidence that belongs together?

↦ **Key takeaway:**
Future priorities focus on live capture, adaptive controls, and signal correlation.

Critical capabilities to address AI-driven identity risks

Future priorities emphasize live capture and adaptive controls, while explainability remains less central.

Live presence

46% Verify biometric capture and human presence

Adaptive controls

34% Trigger additional checks on risk

33% Correlate document, biometric, and session data

32% Assess data integrity and signal source

Governance and consistency

31% Enforce controls across channels

31% Assign accountability for AI decisions

27% Reconstruct and explain decisions

Decision Accountability

Organizations can often trace identity decisions. Fewer can fully explain them.

Not Fully Explainable Decisions

92%

say identity decisions are fully or mostly reconstructable — but only 50% can trace them end to end.

This is an important gap. High-level traceability may show which systems were involved, but not how signals were interpreted or why the final decision was made.

That difference becomes critical when decisions are challenged by regulators, auditors, courts, customers, or internal fraud teams.

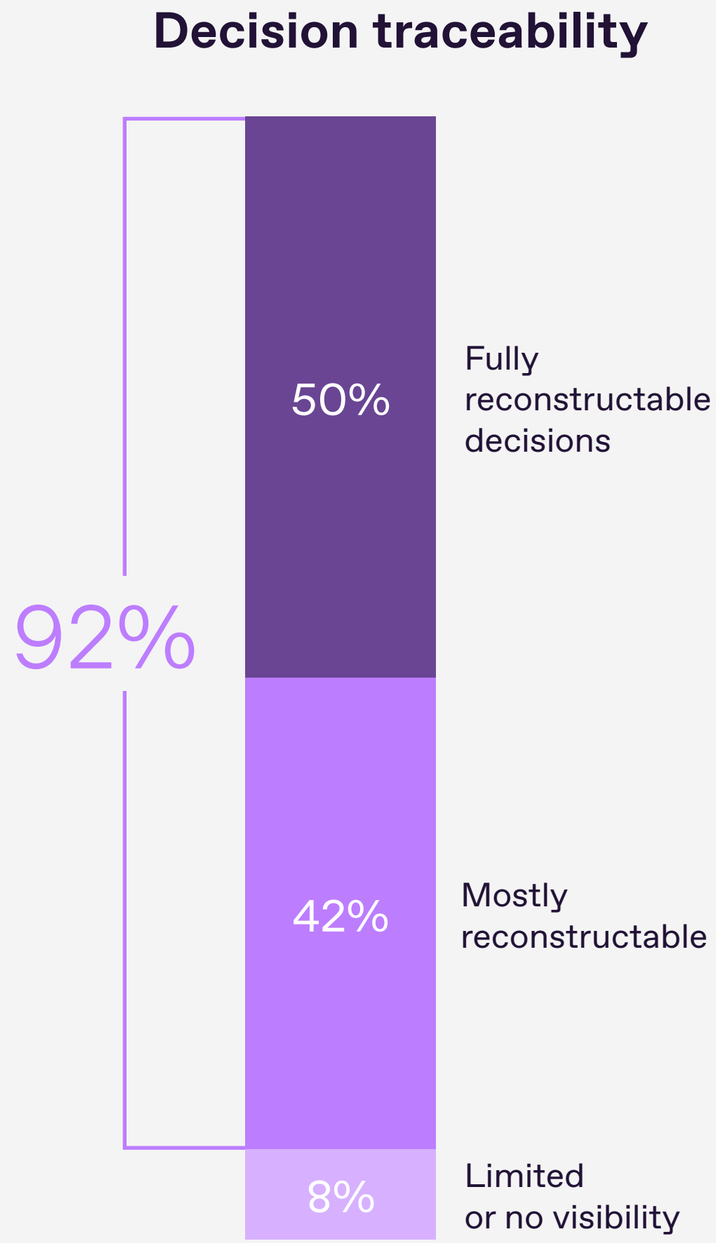

Most organizations can reconstruct decisions to some degree, but only half can explain the full decision chain.

Regulatory Pressure

82% have been required to justify identity decisions externally

56%	26%	15%
Provided audit-grade technical evidence	Provided limited or indirect evidence	Not so far

Regulatory scrutiny experience

External scrutiny is widespread, while evidence quality remains uneven.

This shows that decision accountability is not a future concern. Regulators, auditors, partners, courts, and customers are already asking organizations to explain and defend identity outcomes.

But evidence quality remains uneven. Some organizations can provide audit-grade technical evidence. Others rely on partial logs, vendor reports, or indirect proof.

AI Risks Accountability

AI-driven identity risk does not have a single clear owner.

Ownership of AI-driven identity risk remains fragmented across business, security, fraud, and compliance functions.

This fragmented ownership can slow response. When an AI-assisted identity incident occurs, teams need to know who investigates it, who owns the evidence, who updates controls, and who explains the decision.

Accountability for AI-related identity errors

AI-driven identity risk is entering strategy, but many organizations remain in partial or planning mode.

Shared primary ownership

- **27%** Business leadership
- **27%** Security teams
- **27%** Risk or compliance teams

Residual / outside the core

- **11%** External vendors
- **6%** Accountability is unclear

Adapting to AI

From point checks to connected signals.

Point Check

Connected Signals

How Identity Verification Should Adapt to AI

Identity verification can no longer rely on isolated checks.

Identity verification must move from point-in-time checking to connected assurance. Organizations need to evaluate not only whether a document is valid or a face matches, but also whether the interaction is live, human-controlled, and consistent across the full journey.

This is especially important for AI agents. Automated or AI-assisted actors may combine several signals that look legitimate in isolation.

Connected assurance helps organizations see the full picture: signal source, capture method, data integrity, risk context, and decision evidence.

Connected assurance

AI Visibility as the Maturity Marker

Companies that can see AI activity report stronger controls.

Organizations with clear visibility into AI-assisted identity activity report stronger verification, governance, and evidence capabilities.

This suggests that visibility is not just about detecting AI activity. It reflects whether a company can connect identity signals, understand risk context, and respond across the identity lifecycle.

When AI activity is hard to see, gaps often extend beyond monitoring — into controls, coordination, and decision accountability.

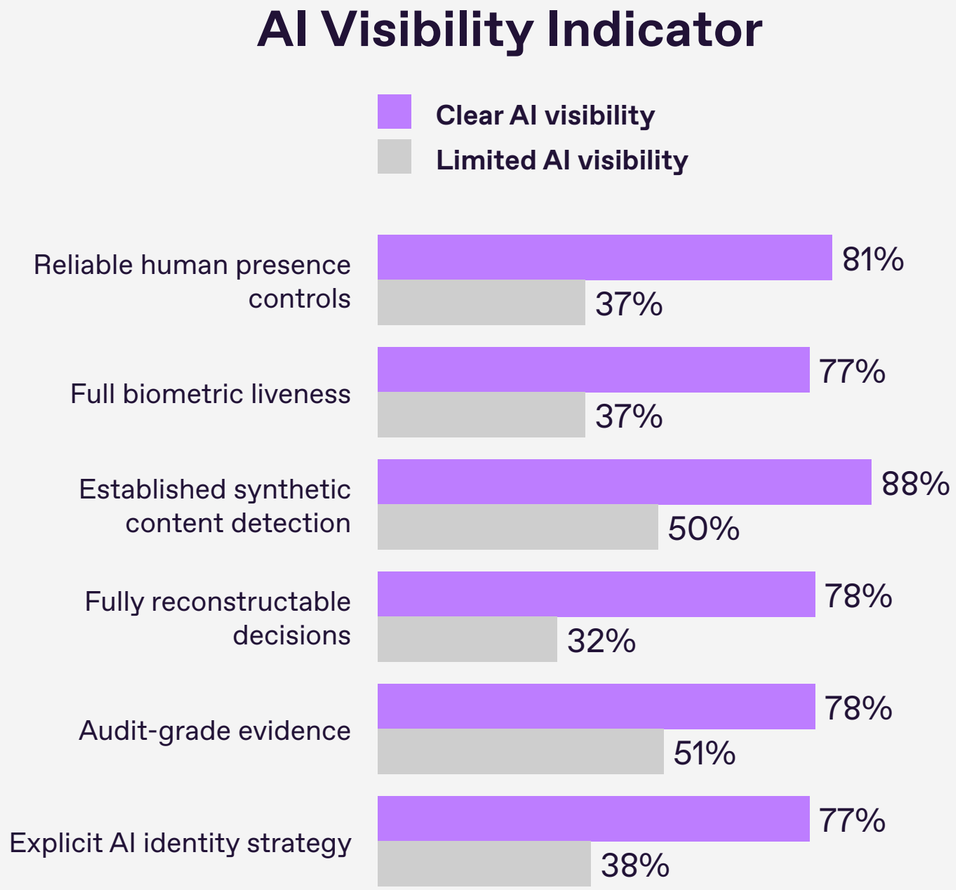

Organizations with clear visibility into AI-assisted identity activity consistently report stronger verification, governance, and evidence capabilities across the identity lifecycle.

Why Is Facial Matching No Longer Enough?

Human presence is becoming the root of trust.

In remote identity verification, a face match is no longer enough. Biometric data can be replayed, injected, synthetic, or deepfake-generated. Organizations must prove that the signal came from a real person, captured live.

That means human presence controls should answer three questions:

1 **Can the system verify live capture?**

2 **Can it detect presentation attacks, injection, and synthetic media?**

3 **Can it connect biometric evidence to the document, session, and risk context?**

AI-driven identity risks make human presence part of identity signal integrity: evidence that the person behind the interaction is real, live, and connected to the identity being verified.

Strategy Is Not Just Paperwork

Organizations that explicitly address AI-driven identity risk in their strategy consistently report stronger verification, governance, and evidence capabilities. This suggests that strategy is not only a policy exercise. It becomes meaningful when it is translated into workflows, escalation rules, evidence requirements, and adaptive controls.

In identity verification, strategy must define what happens when risk changes: when to request stronger proof, when to trigger liveness, when to correlate document and biometric signals, and when to route a case for review.

AI identity strategy as a driver of operational maturity

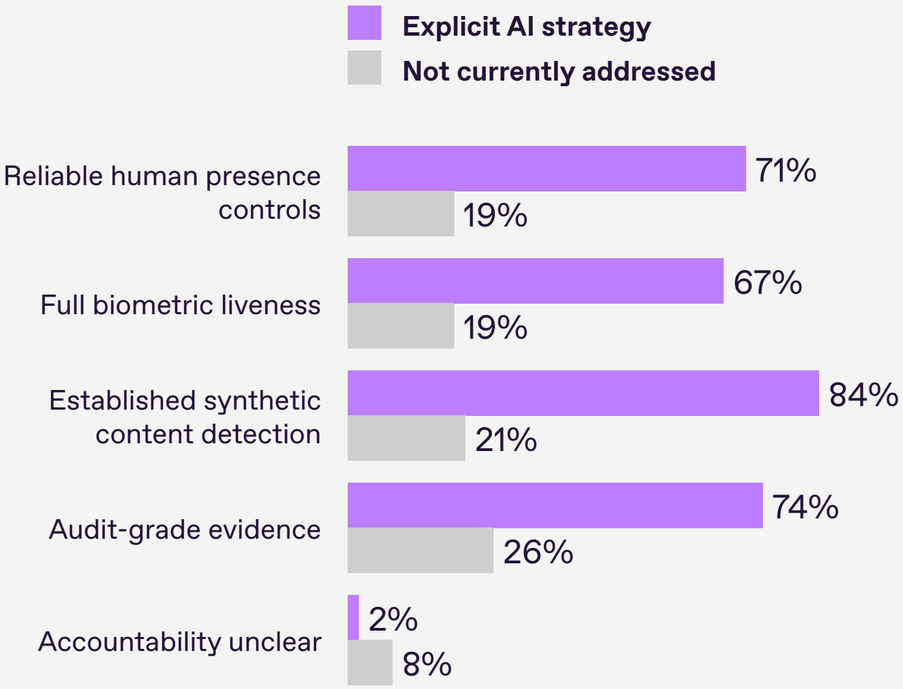

Organizations with explicit AI identity strategies report significantly stronger verification, evidence, and governance capabilities than organizations that do not currently address AI-driven identity risk.

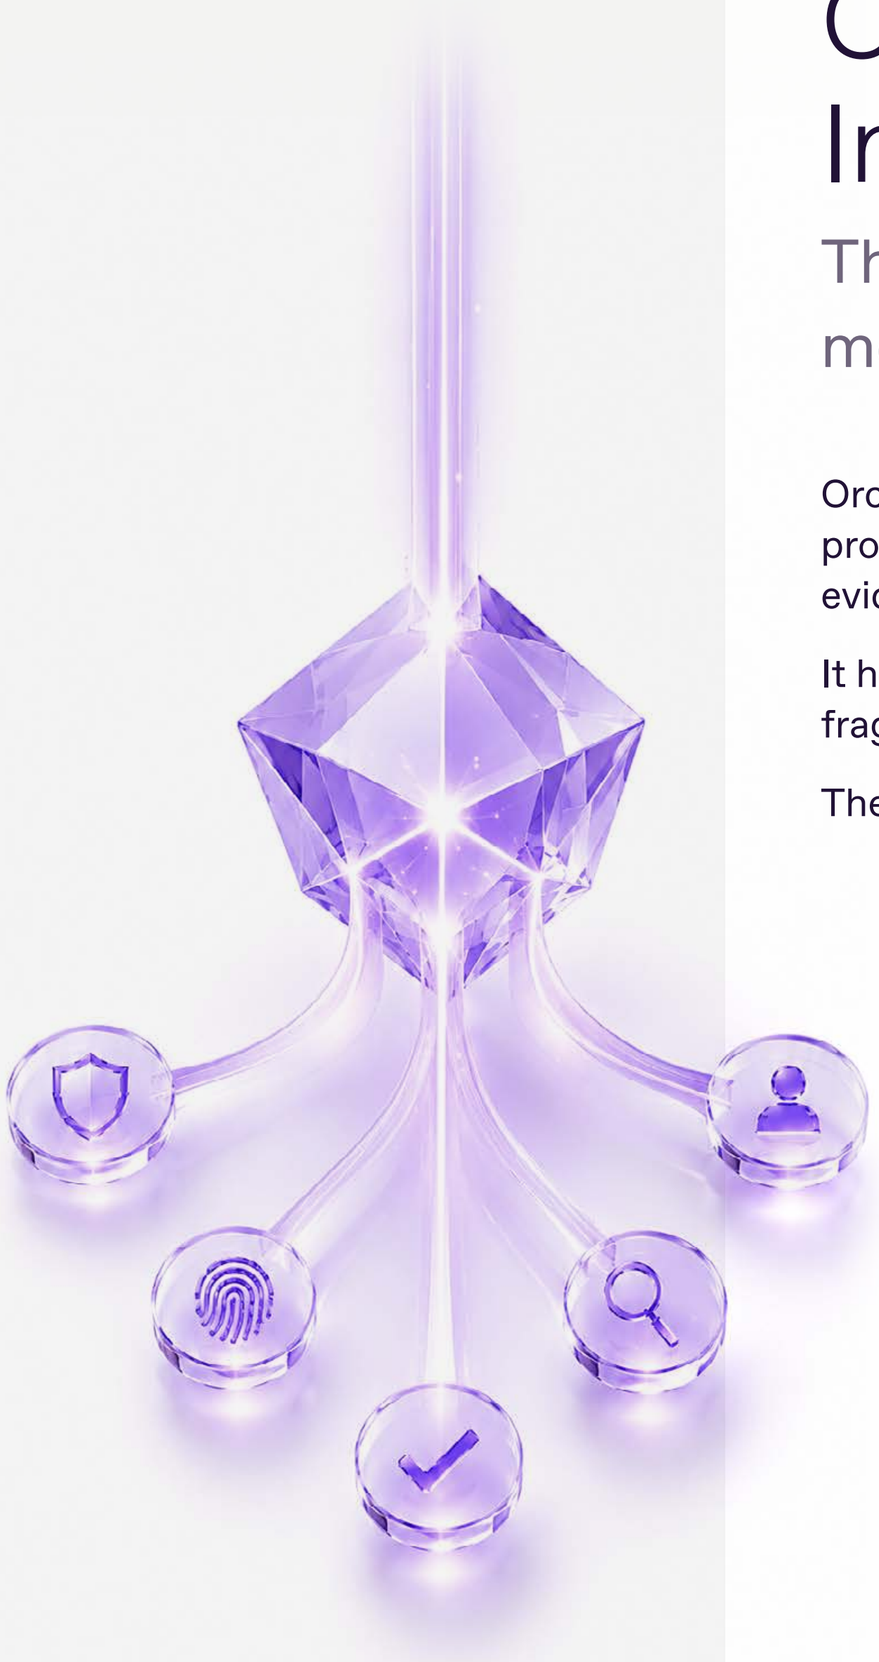

Orchestration Turns Strategy Into Decisions

The right check should happen at the right risk moment.

Orchestration of the identity verification process defines when to request stronger proof, trigger liveness, add AML or PEP screening, route a case to review, and store evidence.

It helps identity checks adapt to risk, geography, regulation, and user type — without fragmenting the decision process.

The goal is not more checks. It is the right assurance level for the right risk.

Decision Reconstructability as an Assurance Indicator

Trusted identity decisions must be explainable later.

Organizations that can fully reconstruct identity decisions report stronger verification, detection, and evidence capabilities. That is because explainable decisions require connected evidence: document checks, biometric results, liveness proof, database screening, session signals, and manual review outcomes tied to one decision record.

When evidence is connected, organizations can show not only what decision was made — but why it was made.

Decision reconstructability as an indicator of identity assurance maturity

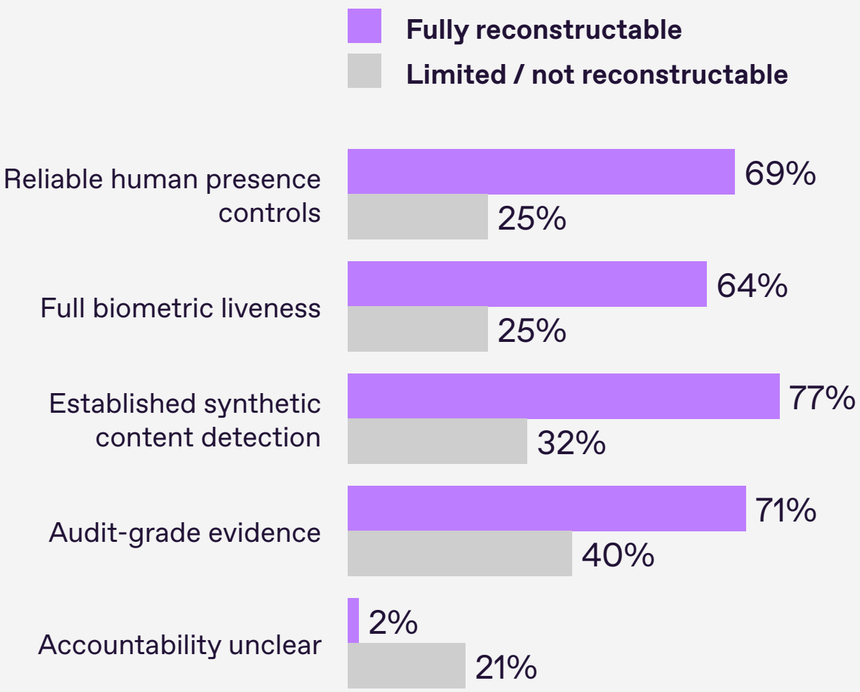

Organizations that can fully reconstruct identity decisions consistently report stronger verification, evidence, and accountability capabilities across the identity lifecycle.

Evidence Must Be Built Into the Workflow

Auditability cannot be added after the fact.

A trusted identity decision is one that can be reconstructed. Organizations must show which signals were used, how they were captured, what checks were triggered, and why the decision was made.

Evidence should be captured as part of the identity workflow, not collected manually, across tools and vendors, after something goes wrong.

Connected evidence turns identity verification from signal checking into decision governance.

Orchestration Strengthens Assurance

AI-driven identity attacks do not stop at one checkpoint.

They can move across onboarding, biometric capture, account access, support, and transaction flows.

Organizations using a platform approach report stronger confidence in verification, evidence, and assurance than the global average.

The implication is clear: identity verification works better when workflows, signals, evidence, and governance are coordinated rather than fragmented.

Platform-based identity architectures and assurance maturity

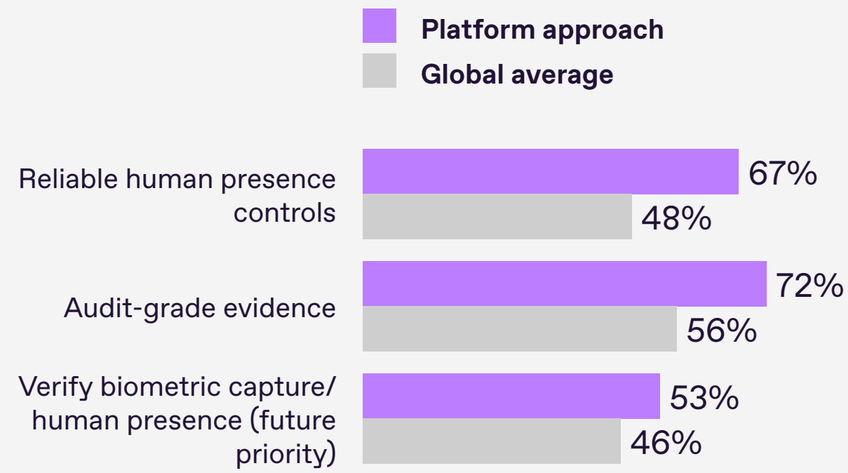

Platform-based IDV environments appear more aligned with assurance requirements.

Why Platform Architecture Matters

Platform architecture is about control, not only efficiency.

AI-driven identity risk moves across the full user journey: onboarding, biometric capture, account access, support, and transactions.

A platform-based approach helps connect document verification, biometrics, liveness, database checks, AML/PEP screening, age assurance, review workflows, and decision evidence in one identity process.

The value is consistency: connected signals, configurable workflows, centralized evidence, and clearer decision context.

Identity verification needs to operate as decision infrastructure — not a collection of separate checks.

The New Identity Architecture

Identity verification is becoming a system for governing trust.

Future-ready identity systems need to answer five questions:

1. **Can the document be trusted?**
2. **Was the biometric captured live?**
3. **Do the signals belong to the same person?**
4. **Does the session context support the identity story?**
5. **Can the decision be reconstructed later?**

That requires identity architecture built around connected signals, adaptive workflows, and preserved evidence.

Identity verification is becoming a system for governing trust across the full customer lifecycle.

Identity Verification Becomes Identity Lifecycle Management

Identity risk does not end after onboarding.

Users return to log in, recover accounts, update data, make transactions, or perform higher-risk actions. Each moment may require a different level of assurance.

That is why identity verification is now a core part of identity lifecycle management.

Organizations need to coordinate document, biometric, database, risk, and review signals across the full customer journey.

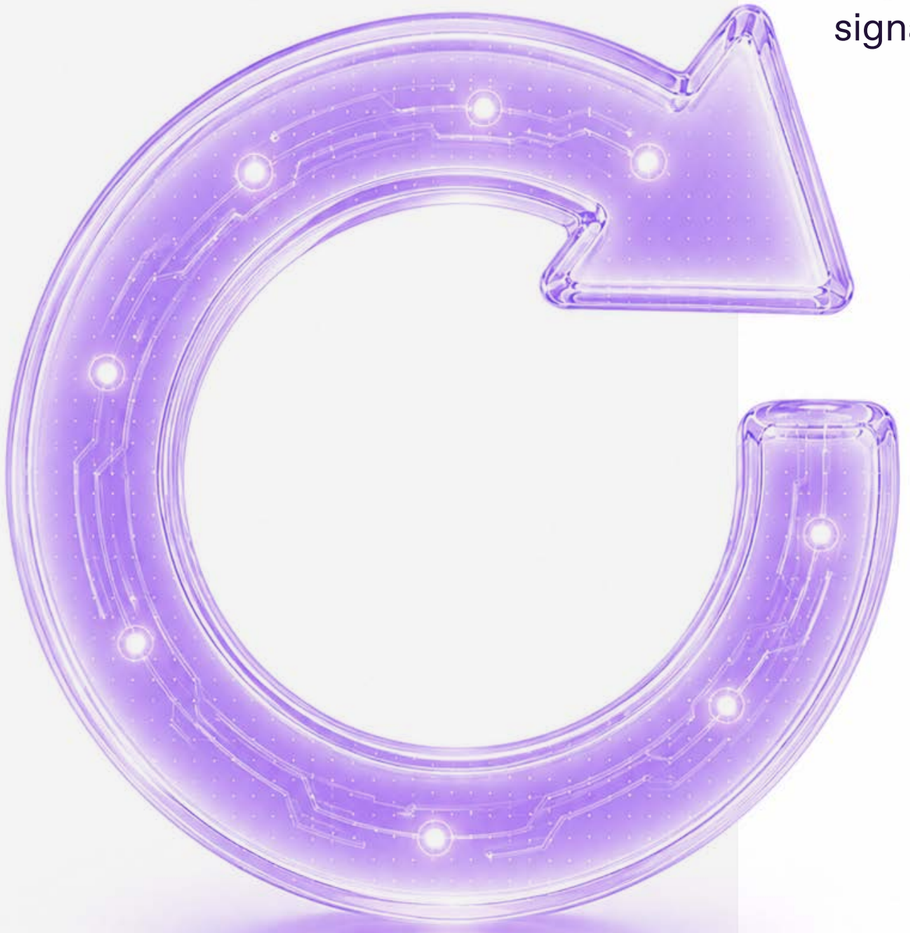

How Regula IDV Platform Addresses Key Pain Points

Regula IDV Platform provides full Identity Lifecycle Management that helps organizations orchestrate identity checks, verify human geniunes, preserve evidence, and manage trust across the full customer lifecycle.

1. **Full identity journey in one platform.** Connect document verification, biometrics, liveness, screening, age assurance, and decision evidence across the full customer lifecycle.

2. **Adaptive workflow orchestration.** Trigger the right checks based on risk, geography, regulation, user type, and transaction value.

3. **Connected identity signals.** Link document, biometric, device, database, and risk signals to detect AI-driven activity that may look legitimate.

4. **Centralized identity management.** Keep identity data, user profiles, verification results, and analytics in one source of truth.

5. **Case management.** Provides a centralized workspace for managing identity, compliance, and fraud investigations, enabling consistent decision-making, efficient collaboration, and end-to-end auditability.

6. **Compliance and auditability by design.** Support KYC, AML, GDPR, CCPA, age assurance, and audit trails to explain identity decisions.

7. **Lower complexity, faster scaling.** Consolidate multiple IDV tools to reduce integrations, simplify operations, and scale faster.

Identity is no longer verified once.

It is maintained across the customer lifecycle.

Practical Takeaways

What the data means for the next 12 months.

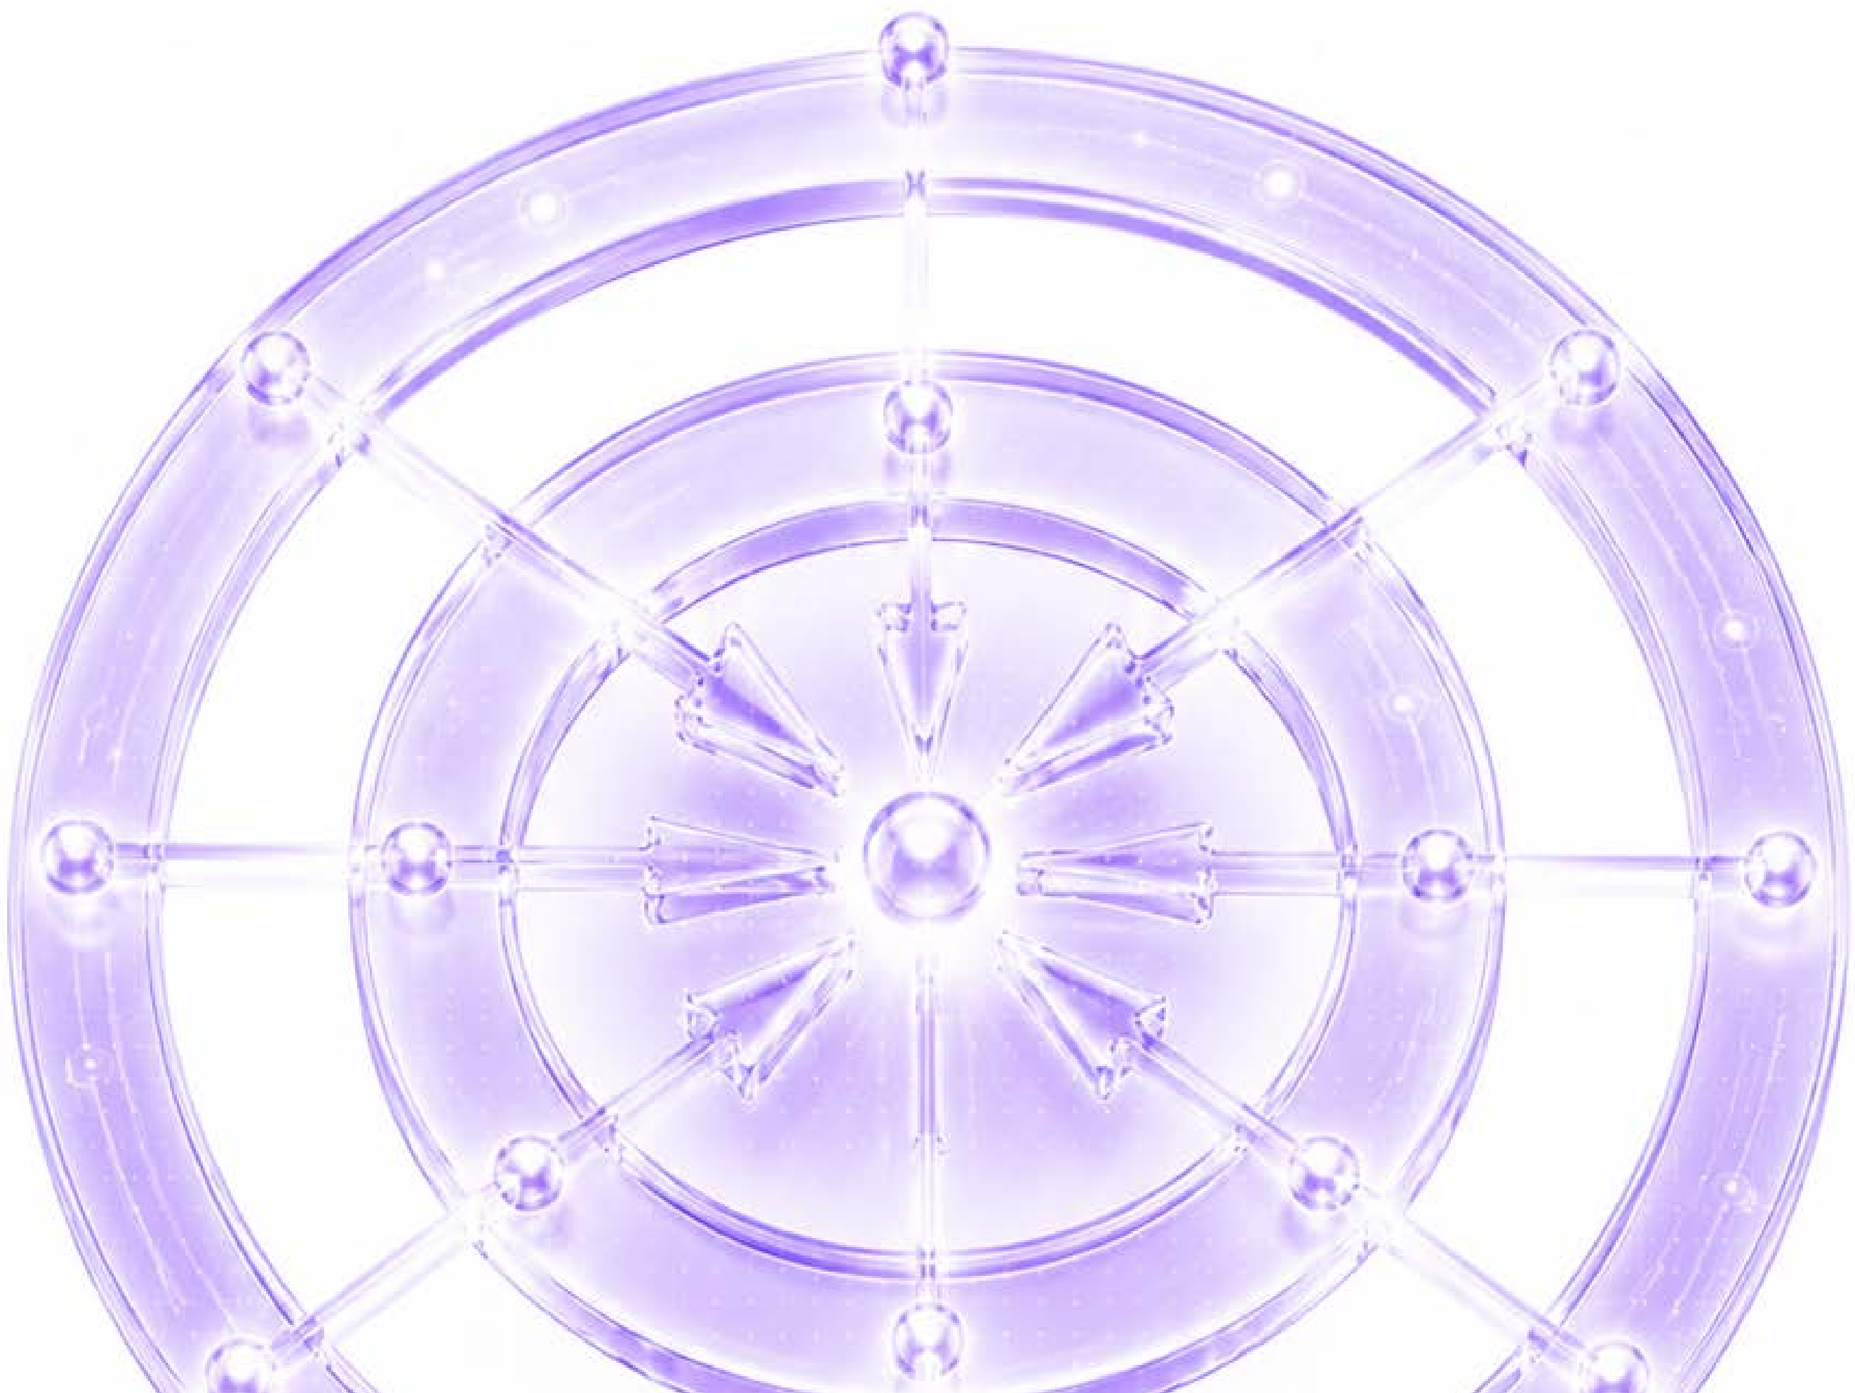

What Organizations Should Do Next

Apply these controls across the full identity lifecycle: onboarding, login, recovery, data updates, and high-risk transactions.

Connect signals. Bring together document, biometric, device, session, risk, and evidence data.

Prove human presence. Use liveness and injection attack detection during biometric capture.

Detect manipulation. Combine document authenticity, biometric integrity, device/session signals, and synthetic content checks.

Adapt workflows to risk. Adjust by geography, regulation, user type, transaction value, and threat signals.

Preserve explainable evidence. Keep logs, signal provenance, check results, and review actions so decisions can be reconstructed.

Outcome:
Identity decisions that are trusted, explainable, and reviewable.

Strengthen Proof of Human Presence

See how liveness detection helps verify that biometric data is captured from a real person in real time.

↗ Contact us: pr@regulaforensics.com
Learn more at regulaforensics.com