



Charting the 2026 Fraudscape

Mid-Year Fraud, Financial Crime, & Cyber Risk Report

Inside this briefing

Five defining market trends

The strategic implications for risk leaders

A practical playbook for navigating the second half of 2026

Foreword

Navigating by One Star

For most of the last decade, fraud, financial crime, and cyber risk were treated as three separate problems, owned by three separate teams, watched through three separate screens.

The first half of 2026 has made it clear that **the adversary no longer respects those boundaries**, if it ever did. A single attack now begins with an AI-generated lure, runs through a compromised or synthetic identity, executes as an instant payment, and settles through a stablecoin, crossing the lines between cyber, fraud, and money laundering in a matter of seconds.

This shift makes the cost of a separated approach impossible to ignore.

This briefing is Realta's read on where the threat is heading at the half-year mark.

Our objective is not simply to describe the market but to help risk, fraud, payments, and compliance leaders prepare for the operational decisions the second half of 2026 will demand.

It is written for the people who have to make the call, in real time, on whether to approve, decline, or review, and who are measured equally on the fraud they stop and the good customers and legitimate revenue they keep.

Throughout, we have tried to separate signal from noise: to report what the data actually shows, to be honest about what is still uncertain, and to point to the operational implications rather than the headlines.

Our own conviction, reflected in how we built Realta, runs through these pages: that risk, fraud, and compliance are most effective when they share one view of the customer, one stream of signals, and one explainable decision.

Risk, fraud, and compliance are most effective
when they share one view of the customer, one
stream of signals, and one explainable decision.

We have flagged that perspective openly where it appears, marked as Realta's view, so you can take the analysis and the opinion separately. The trends, however, belong to everyone working in this field.

Our hope is that this briefing serves as a reliable point of reference wherever you sit—a clear star to guide your team through rapidly changing threats.

Executive Summary

The first half of 2026 confirmed that **fraud and financial crime have become a single, AI-accelerated, industrialized economy**, and that the migration to real-time payments has handed that economy a structural advantage: a settlement window now measured in seconds.

For organizations across the payment lifecycle, from onboarding a merchant to monitoring a transaction to filing a suspicious activity report, the defining challenge is no longer detecting any one threat.

It is seeing the whole picture fast enough to act.

Five findings that define the half-year

AI is now on both sides of the transaction.

01

Industry surveys attribute more than half of fraud to AI-assisted methods; deepfakes, voice cloning, and generative phishing are now baseline tradecraft, and roughly nine in ten institutions are deploying AI in defense. It takes AI to counter AI.

Scams have scaled like an industry.

02

Authorized push payment and investment scams are run from industrial compounds; fewer victims are hit for far larger sums; and the real-time rails that delight customers also make stolen funds irrecoverable in seconds.

The merchant edge is widening.

03

Card-not-present fraud, chargeback and refund abuse, account takeover, digital-wallet risk, and the first wave of agentic-commerce fraud are blurring the line between criminal and customer.

Crypto, and stablecoins above all, is the laundering layer.

04

Record illicit on-chain value in 2025, the great majority of it in stablecoins, is where fraud proceeds, sanctioned flows, and cartel money now converge.

Regulation is fragmenting on three fronts.

05

Europe is centralizing AML supervision while the U.S. deregulates; a new front of AI regulation is opening across 70-plus jurisdictions; and a growing group of countries is making banks, telcos, and platforms share liability for scams.

Strategic Implications

Customer friction is becoming a competitive disadvantage. Treating fraud as an unavoidable cost of doing business, and tolerating the customer friction that legacy controls create, is itself a competitive liability.

Convergence is the winning posture. Institutions that connect identity, behavior, transaction, and compliance signals into one explainable decision will outperform those still making risk decisions in separate systems.




Realta's View

Every theme in this briefing points the same way: the adversary has organized to exploit the seams between your systems.


The only durable answer is to remove the seams to deliver one customer view, one signal stream, and one explainable decision in real time across the entire payment lifecycle.

The Numbers Defining 2026



U.S. internet-crime losses have nearly quintupled since 2020

Reported losses to the FBI Internet Crime Complaint Center (IC3), 2020-2025



Source: FBI IC3 Annual Internet Crime Report

Behind the totals sit measurable operational pressures. In the U.S. FBI’s 2025 Internet Crime Report, **complaint volumes and losses both rose sharply year over year**, and for the first time the Bureau broke out AI-related complaints as its own category—with more than 22,000 of them accounting for roughly \$893 million in reported losses.

In Europe, the European Banking Authority recorded payment-fraud losses climbing to around €4.2 billion. And in the Federal Reserve's 2026 Risk Officer Report, a clear **majority of institutions reported rising fraud across every payment channel**, with impersonation, social engineering, and credential compromise the common threads.

Three forces explain the step change, and they structure the rest of this briefing: the arrival of AI on both sides of the transaction, the industrialization of scams against the backdrop of real-time payments, and the maturation of crypto as a laundering layer, all unfolding against a regulatory map that has begun to split.

NB: Where this briefing cites full-year totals, they are drawn from the annual datasets published across the sector in the first quarter of 2026; the trend, enforcement, and regulatory developments are those observed from January through early June. Figures from different sources use different methodologies and are not always directly comparable; they are presented to show scale and direction, not as a single reconciled dataset.

TREND 01

AI Is Now on Both Sides of the Transaction

AI has removed the traditional cost, skill, and scale barriers to sophisticated fraud, forcing institutions to rethink how trust is established and maintained.

The defining feature of the 2026 fraudscape is the operational maturity of AI, for attacker and defender alike.

In the World Economic Forum’s Global Cybersecurity Outlook 2026, a survey of more than 800 senior leaders across 92 countries, **94% named AI the single most significant driver of change in cybersecurity, and 87% called AI-related vulnerabilities the fastest-growing risk of the past year.** Perhaps more strikingly, for the first time, chief executives ranked cyber-enabled fraud as their top concern, ahead of ransomware.

The reason is structural: generative tools that produce convincing text, voice, and video are now free or near-free, require no technical skill, and can be used anonymously and at scale. This combination of zero cost, zero skill, and zero accountability has made AI-enabled fraud the fastest-growing category.

The Attacker’s Toolkit


Deepfakes and synthetic media	Synthetic media now features in roughly one in nine fraud attempts globally, and identity-verification vendors report around one in five biometric fraud attempts involving a deepfake. The deepfaked video call that cost engineering firm Arup roughly \$25 million has become the reference point for executive-impersonation fraud.
Voice cloning	A few seconds of audio is now enough to clone a voice, powering the “distressed relative” and payment-instruction scams that fall hardest on older and more vulnerable customers.
Generative phishing at scale	AI-written lures achieve markedly higher response rates than human-crafted ones, while a single operator can run thousands of personalized, context-aware conversations at once.

Synthetic identity

Fabricated but plausible customers remain among the hardest fraud to detect because the applicant looks legitimate at onboarding; in the UK, Experian reported AI-related fraud rising from roughly 23% of cases in 2024 to about 35% in early 2025.

Generative-AI-enabled fraud in the U.S is projected to triple

Estimated GenAI-enabled fraud losses



Source: Deloitte Center for Financial Services projection

The practical consequence is that identity can no longer be verified once, at the door, with a document scan or a selfie. The methods that anchored digital trust are being defeated, and **verification must become continuous, behavioral, and multi-signal across the customer lifecycle.**

The defender's response

The encouraging half of the story is that institutions are not standing still. In the same survey, around nine in ten institutions report using AI to detect fraud and accelerate investigations, applying it across scam detection, transaction fraud, and anti-money-laundering work.

The lesson the sector has internalized is blunt: rules alone cannot keep pace, and it takes AI to counter AI.

Adoption has surfaced two hard constraints.

The first is data: practitioners consistently name fragmented, siloed, inconsistent data that is spread across card, online, and mobile channels as the single biggest barrier to effective AI, because a model is only as good as the unified view it is trained and run on.

The second is trust: as models make more consequential decisions, the demand for transparent, explainable, auditable reasoning, rather than a black box, has moved from a nice-to-have to a regulatory and commercial necessity.

Realta's View

The data barrier is the real barrier. **AI does not fail for lack of models; it fails for lack of a unified, low-latency view of the customer and the transaction.**

A shared feature store that turns raw onboarding and transaction streams into real-time signals, and decisions that can be explained rather than merely scored, are what separate AI that works in production from AI that works in a slide. That is why Realta has built a bespoke feature store tailored to the world of fraud and financial crime.

TREND 02

The Scam Economy and the Real-Time Squeeze

Real-time payments have transformed scams from a fraud problem into a race against the clock.

Authorized push payment fraud, in which a victim is manipulated into willingly sending money, has become **the defining consumer-facing financial crime of the era**, and it is the clearest example of how cyber, fraud, and organized crime have fused.

The schemes are seeded online and over telecoms networks, executed through social engineering, and monetized through laundering networks, with no single perimeter to defend. Modeling by ACI Worldwide projects APP losses across the U.S., U.K., and India climbing toward roughly \$5.25 billion by 2026, a compound growth rate above 20%, with product, romance, and investment scams the dominant subtypes.


...APP losses across the U.S., U.K., and India climbing toward roughly \$5.25 billion by 2026...

Fewer victims, larger losses

The pattern across the period is consistent. In the UK, APP values rose even as case volumes fell in some categories, with investment scams a standout. Roughly three-quarters of UK cases originate online, typically lower-value purchase scams, while the smaller share originating via telecoms, principally impersonation, accounts for a disproportionate share of total losses. The mandatory reimbursement regime introduced by the UK Payment Systems Regulator, live since October 2024, now splits the cost of reimbursing most victims between sending and receiving institutions up to a per-claim cap of £85,000, permanently changing the economics: **liability now sits with the banks on both sides of a fraudulent payment.**

Fewer victims, larger losses

UK authorized push payment (APP) fraud, by scam origination channel



Source: UK Finance. Online cases are higher-volume/lower-value; telecoms cases are lower-volume/higher-value.

Industrial scam compounds

The supply side has industrialized. Investigators describe vast scam compounds across Cambodia, Myanmar, and Laos, run largely by transnational criminal organizations and staffed in significant part by trafficked, coerced labor. INTERPOL has warned the model has gone global. Enforcement scaled in response: a U.S. Department of Justice Scam Center Strike Force seized more than \$578 million in cryptocurrency in its first three months, and the FBI's proactive notification effort had alerted nearly 9,000 investment-fraud victims by early 2026, most of whom did not yet realize they were being scammed.

The real-time squeeze

Underpinning all of this is the global migration to instant payments. Faster settlement improves customer experience but compresses the window for detection and intervention to seconds, and fraudsters are deliberately targeting instant rails precisely because funds become irrecoverable so quickly. **This is the operational seam where the fraud and cyber disciplines must merge:** an account takeover is a cybersecurity event that produces a fraud loss, and stopping it in real

time requires behavioral, device, and transaction signals working together before the payment leaves.

 Realta's View

Reimbursement has turned scams into a balance-sheet problem for both banks in a transaction, meaning that not just outbound-payment risk but also inbound has to be scored in real time. **When the settlement window is seconds, the decision window has to be under one second.**

Pre-transaction intervention, name-checking, mule-account detection, and behavioral analytics on both sides of the payment are no longer optional.

TREND 03

Fraud at the Merchant Edge

Merchant risk begins long before the transaction and extends far beyond checkout.

For payment processors, acquirers, marketplaces, and the merchants they serve, the half-year’s most important shift is the blurring of the line between the criminal and the customer. Merchant surveys from across the ecosystem describe **fraud that is becoming more diffuse, more first-party, and harder to separate from legitimate behavior**, even as outright criminal attacks grow more automated.

Card-not-present and chargebacks	CNP fraud and its disputes remain the baseline cost of online commerce. The friction vs. conversion trade-off is sharper than ever: over-tightening loses good revenue, loosening invites loss.
Refund and policy abuse	First-party and "friendly" fraud, including returns abuse, is now endemic. Abusers succeed at very high rates, and a small core of repeat offenders drives disproportionate loss, forcing organizations to try to improve fraud protection without punishing good customers.
Account takeover and multi-accounting	Credential stuffing, bot-driven login attacks, and multi-account promotion abuse sit at the seam between security and fraud.
Digital-wallet risk	Merchants increasingly rank mobile wallets among their most fraud-prone payment methods. Liability is different than in CNP flows and can catch teams off guard if they are accustomed to chargeback rules.
Agentic commerce	AI shopping agents create a largely untested fraud surface with unresolved liability. Whatever the adoption rate of agentic commerce, its fraud is coming. Organizations must plan for it now.

Onboarding and supplier fraud

For platforms and marketplaces, risk extends to onboarded sellers and suppliers; synthetic businesses entering at underwriting become every downstream party's problem.

The cost of a false decline is now a strategic business risk.

The common thread is that the **merchant edge can no longer be defended by a single control at the checkout**. Risk now has to be assessed at onboarding, continuously through the account lifecycle, and at the moment of each transaction, with the same view of identity and behavior carried across all three.

Risk across the payment lifecycle



powered by one decisioning engine



Realta's View

Acceptance and protection are not opposing goals; they are the same goal seen from two sides. The objective is to approve more good business while blocking more bad, which is **only possible when underwriting, transaction monitoring, and dispute signals feed one model** of the customer rather than three disconnected ones.

A decline that should have been an approval is as much a failure as a fraud that should have been a decline.


TREND 04

Crypto and Stablecoins: The Laundering Layer Matures

Stablecoins have become the preferred infrastructure for moving illicit value at global scale.

If AI is the production engine of modern fraud and if scam compounds are the factory floor, **cryptocurrency—and stablecoins above all—have become the logistics network that moves the proceeds.** Chainalysis estimates illicit cryptocurrency received by identified addresses reached at least \$154 billion in 2025, a 162% year-over-year increase, with growth driven heavily by activity tied to sanctioned entities. Critically, stablecoins accounted for roughly 84% of that illicit volume, a decisive shift toward dollar-pegged tokens that are easier to transact and value-stable.

Stablecoins now carry the great majority of illicit crypto value



Source: Chainalysis 2026 Crypto Crime Report. Share of 2025 illicit crypto volume.

Laundering tradecraft has professionalized accordingly: networks hold value briefly in less freeze-prone assets before converting into stablecoins near the cash-out point and then move funds through decentralized exchanges and cross-chain bridges to frustrate tracing. The average value sitting in scam-linked wallets has fallen even as volumes rose, evidence of faster turnover. A meaningful counterweight is the growing cooperation of stablecoin issuers with law


enforcement; several 2026 seizures were enabled by issuers freezing and transferring assets, a control point that does not exist for fully decentralized assets.

The essential point for financial-crime teams is convergence: **the same on-chain infrastructure now carries scam proceeds, sanctioned-entity flows, narcotics money, and state-sponsored theft**, with North Korea-linked actors reported to have stolen more than \$2 billion in 2025. Moreover, exposure to any one typology increasingly means exposure to all of them, because they share rails.

Blockchain analytics has moved from a crypto-native niche to a core control for any institution touching faster payments, fintech partnerships, or stablecoin settlement.

The average scam payment more than tripled in a year

Average value per scam payment



Source: Chainanalysis 2026 Crypto Crime Report.



Realta's View

Stablecoins are becoming part of mainstream payments, which means stablecoin-aware screening is becoming part of mainstream payments risk. The institutions that fare best will treat on-chain visibility as one more signal in a unified decision, **not as a separate, bolted-on system that only the crypto desk ever sees.**

TREND 05

A Fragmenting Regulatory Map: AML, AI, and the Cross-Sector Turn

Regulatory expectations are diverging across jurisdictions while converging around one requirement: explainable decisions.

For most of the last decade, the direction of financial-crime regulation was broadly convergent: more obligations, more transparency, more supervision.

The first half of 2026 made it clear that **consensus has fractured, and on three fronts at once**. Anti-money-laundering regimes are diverging between a centralizing Europe and a deregulating United States; an entirely new front of artificial-intelligence regulation is opening across more than seventy jurisdictions; and a growing group of countries is rewriting who is liable for scams.

Institutions that operate across borders must now reconcile several genuinely different philosophies at the same time.

Anti-money laundering: Europe centralizes, the U.S. deregulates

The EU is consolidating two decades of fragmented rules into a single regime. Its new Authority for Anti-Money Laundering (AMLA) will directly supervise the highest-risk cross-border institutions as well as certain non-EU firms operating in the bloc.

*This is the year to prepare, as the AMLA is
building out its supervisory machinery...*

This is the year to prepare, as the AMLA is building out its supervisory machinery and delivering roughly two dozen technical standards in July 2026. The centrepiece is the directly applicable AML Regulation, the "single rulebook" imposing uniform requirements on due diligence, beneficial ownership, and sanctions from 10 July 2027.

The US moved the opposite way: in 2026, FinCEN eased beneficial-owner reverification, proposed refocusing the Bank Secrecy Act from checklist compliance to risk and outcome evaluation, and dramatically narrowed the Corporate Transparency Act reporting to foreign companies. But enforcement escalated to over a thousand OFAC designations targeting cartels and scam networks. One sanction of a cartel cash-to-crypto network included specific crypto addresses, putting every virtual asset provider touching them on notice.

Dimension	European Union	United States
Direction of travel	Centralizing, expanding obligations	Deregulating program & reporting burden
Supervision	New single authority (AMLA)	Risk-based, outcome-oriented BSA reform
Beneficial ownership	Single rulebook, broad scope from 2027	CTA narrowed; >99% of entities exempt
Where pressure is rising	Supervisory consistency, non-EU reach	Sanctions & fraud enforcement (cartels, crypto)

Artificial intelligence: a new regulatory front


Layered on top of the AML divergence is an entirely new front. By early 2026, more than seventy countries had launched over a thousand AI-policy initiatives, and, like AML, they are splitting rather than converging.

...an entirely new front of artificial-intelligence regulation is opening across more than seventy jurisdictions...

This matters directly for financial crime because **the systems that now make consequential decisions**, credit scoring, identity verification, fraud and transaction risk, **are precisely the AI systems regulators have moved to govern**, with new

expectations around documentation, bias testing, human oversight, and the ability to explain an automated decision.

AI regulation is splitting, not converging



Illustrative positioning of major jurisdictions by regulatory intensity for AI 2026.

The EU has the most developed framework.

The AI Act took effect in 2024 on a staggered schedule, with prohibitions applying from February 2025 and general-purpose-model rules from August 2025. Its high-risk regime covering credit scoring and insurance pricing, though not fraud detection, was deferred by the May 2026 "Digital Omnibus" to December 2027. The substance is unchanged: conformity assessments, documentation, bias monitoring, and human oversight, with penalties up to €35 million or 7% of global turnover, applying extraterritorially.

Elsewhere, the US has pivoted to federal deregulation while state laws take effect in 2026, creating a patchwork. China enforces binding sector rules and content labeling; South Korea's 2026 Act also reaches foreign providers; Japan, the UK, and Canada take lighter or still-pending approaches.

Jurisdiction	Approach	What it means for financial-crime AI
European Union	Comprehensive risk-based law (AI Act)	Credit scoring & insurance pricing are high-risk; documentation, oversight, explainability
United States	Federal deregulation vs. active state laws	Patchwork across states; federal push to preempt; persistent uncertainty

United Kingdom	Principles-based, regulator-led	No standalone AI law yet; existing regulators apply current rules
China	Binding, sector-specific rules	Generative-AI controls and mandatory synthetic content labeling
Canada/ Japan / Korea	Mixed and emerging	Canada's AIDA pending; Japan's AI act is principles-based; Korea's is extraterritorial.

For fraud and financial-crime teams **the throughline is explainability**. Whether the test is the EU AI Act’s documentation duties, longstanding rights around automated decision-making, or a regulator simply asking why a customer was declined, the era of the unexplained model score is closing.


A decision that cannot be reasoned about cannot be defended.

The cross-sector turn: sharing the burden for scams

A third shift redraws liability for scams.

The UK's mandatory reimbursement regime splits the cost of most authorized-push-payment fraud between sending and receiving banks. Australia goes further: its Scams Prevention Framework imposes statutory duties on banks, telecoms, and digital platforms, with penalties up to A\$50 million per breach and first obligations from July 2026, redress hinging on "reasonable steps" rather than the loss itself. Singapore similarly splits liability between banks and telcos. The direction is unmistakable: **obligations are moving toward where scams originate, demanding real-time detection and cross-sector intelligence sharing**.

A global regulatory calendar: three fronts, 2025-2027



Realta's View

Fragmentation on three fronts **rewards two capabilities above all.**

The first is configurability, the ability to tune rules, models, thresholds, and reporting per jurisdiction without re-engineering the core, so one platform can satisfy a prescriptive EU regime, a discretionary U.S. one, and a cross-sector scam duty at once.

The second is explainability. The same audit-ready, no-black box reasoning the EU AI Act now demands is also what lets a firm show it took "reasonable steps" against a scam and justify why it declined a customer.

Even as the rules across AML, AI, and scam regulation diverge, the regulator is converging on **a single question: *can you explain this decision?***

THE IMPLICATION

The Convergence Imperative

The institutions that connect fraud, cyber, and financial crime will outperform those that continue managing them separately.

The unifying lesson of the half-year is that adversaries operate across the boundaries that institutions draw between fraud, cyber, and anti-money-laundering, **knowing that a signal split across three siloed systems is a signal nobody acts on**. A scam victim's outbound payment (fraud), the account takeover that enabled it (cyber), and the mule network that received it (AML) are three views of one event.

The institutions that fared best in the first half of 2026 were those beginning to dissolve those internal boundaries, in their data, their teams, and their tooling.

Convergence: three signal streams, one explainable decision



Two practical failures recur.

The first is the data silo: information fragmented across card, online, mobile, onboarding, and monitoring produces a fragmented view of the customer, and no model can outperform the data it sees.

The second is the false-positive tax: legacy, rules-bound controls generate noise that buries real threats, burns analyst time, and adds friction that quietly costs good customers and legitimate revenue. Treating that noise, and the fraud losses alongside it, as an unavoidable cost of doing business is **no longer defensible when better is achievable**.

A converged posture rests on four critical capabilities, each of which maps to a failure the half-year has exposed:

One view. A unified, real-time picture of identity, device, behavior, and transaction, shared across onboarding, transaction monitoring, and AML, so nothing slips through the gaps between systems.

Real-time decisioning. Sub-second approve, decline, or review decisions, because in an instant-payments world the decision window is the settlement window.

One view.

Explainability. Every decision is auditable and understandable, not a black box, to satisfy both regulators and the commercial need to know why good customers are being declined.

One decision.

Operational leverage. Automating the automatable and cutting false positives, so teams scale with risk rather than headcount.

Realta's View

This is the conviction Realta was built on: bring risk, fraud, and compliance into one system, score events in real time against a shared feature store, decide in milliseconds, and surface a clear reason for every call.

Convergence is not a slogan; it is the only architecture that matches how the adversary already operates.

The trends in this briefing point toward a practical operating model. The following playbook is intended to identify priorities leaders can evaluate over the next 12–18 months rather than a one-size-fits-all implementation plan.

THE RESPONSE

A Playbook Across the Payment Lifecycle

This tool is organized around the lifecycle of a payment, from the moment a customer or merchant is onboarded to the moment a transaction is decided to the longer arc of monitoring for financial crime. It is offered as a strategic checklist rather than a product specification; sequencing depends on each organization's starting point and risk profile.

At onboarding and underwriting

Treat identity as continuous, not a one-time gate.

-  Assume documents, static selfies, and single-moment liveness checks can be defeated by synthetic media. Layer behavioral, device, and multi-signal verification that persists beyond account opening.

Screen the business, not just the individual.

-  For platforms and marketplaces, extend diligence to merchants, sellers, and suppliers, where synthetic and misrepresented entities enter the ecosystem.

At the transaction

Close the real-time gap.

-  Invest in sub-second pre-transaction intervention, name-checking, and inbound-payment risk scoring, as this is the operational front line of the reimbursement era and the instant-payments squeeze.

Score behavior, not just transactions.

-  Adopt behavioral analytics that model genuine customer behavior so anomalies, account takeover, and scam-driven payments surface without drowning teams in false positives.

Balance acceptance and protection deliberately.

-  Measure the revenue lost to over-tight controls alongside fraud losses, and tune for both; a wrongful decline is also a failure.

Across financial crime and compliance

Unify the data and the teams.

-  Bring fraud, cyber, and AML onto a shared view and shared case management; technology cannot compensate for organizational silos.

-  **Acquire on-chain visibility.**
Fold stablecoin-aware screening and blockchain analytics into the core control set, not a separate system seen only by a crypto desk.
-  **Engineer for regulatory divergence.**
Run a configurable model that can satisfy an expanding European regime and a discretionary U.S. regime in parallel, with sanctions screening as a live control.
-  **Demand explainability.**
Require that every automated decision can be understood and audited; it is now both a regulatory expectation and a commercial safeguard.

Outlook for the Second Half of 2026

We expect some trajectories set in the first half of the year to intensify, not reverse, through the remainder of 2026.

Agentic AI on both sides. Autonomous agents will automate not just content but the orchestration of multi-step scams and the orchestration of defense, accelerating the arms race. Agentic-commerce fraud in particular will move from emerging to material.

Cross-sector liability pressure. Expect continued pressure to extend fraud-prevention duties to the social-media and telecoms platforms where scams originate—beyond the banks that currently bear the reimbursement cost.

Real-time everywhere. As account-to-account and instant rails spread in commerce, the sub-second decision becomes the default expectation rather than the exception.

AI regulation begins to bite. Even with the EU's high-risk deadline deferred to late 2027, conformity, documentation, and explainability work starts now. U.S. state laws take effect through 2026, and more jurisdictions follow, making explainable, well-documented models a cross-border requirement.

Sustained sanctions escalation. Cartel- and scam-network designations naming crypto addresses will keep widening screening obligations even as other U.S. burdens ease.

Explainability as table stakes. Regulatory and commercial pressure will continue to push opaque models toward transparent, auditable decisioning.

The institutions that treat the first half of 2026 as a turning point and reorganize around convergence rather than incremental tooling will be the ones ahead of their adversaries at year end.

The institutions that see the whole picture first and can explain every decision they make will define the next generation of risk management.

Methodology & Sources

This briefing synthesizes publicly available industry research, regulatory and enforcement disclosures, and threat-intelligence reporting tracked by Realta through the first half of 2026. Full year totals draw on annual datasets published across the sector in the first quarter of 2026; enforcement, regulatory, and threat developments are those disclosed between January and early June 2026. Figures from different sources use different methodologies and reporting scopes and are not always directly comparable.

Selected sources referenced:

- Nasdaq Verafin — 2026 Global Financial Crime Report
- World Economic Forum — Global Cybersecurity Outlook 2026 (with Accenture)
- Chainalysis and TRM Labs — 2026 Crypto Crime Reports (scams, sanctions, stablecoins)
- FBI Internet Crime Complaint Center (IC3) — 2025 Internet Crime Report
- Feedzai — AI Trends in Fraud and Financial Crime Prevention
- Ravelin — Global Fraud Trends and Global Payments reports (merchant surveys)
- Featurespace — State of Fraud and Financial Crime benchmarks
- ACI Worldwide & GlobalData — Scamscope APP fraud projections
- UK Finance and the UK Payment Systems Regulator — APP fraud data and reimbursement framework
- European Banking Authority — payment fraud data; EU AMLA and the EU AML package (AMLR, AMLD6, AMLA Regulation)
- EU AI Act and the 2026 “Digital Omnibus”; U.S. state AI laws; China, South Korea, Japan, Canada and UK AI frameworks
- Australia Scams Prevention Framework Act 2025; Singapore Shared Responsibility Framework; UK Payment Systems Regulator
- U.S. Treasury / FinCEN and OFAC; U.S. Department of Justice; Federal Reserve Financial Services; U.S. GAO

– Global Anti-Scam Alliance (GASA); Deloitte Center for Financial Services; World Economic Forum; Experian

This document is provided for general information and does not constitute legal, regulatory, or compliance advice; organizations should assess their own obligations and risk exposure with appropriate professional support. Realta's view denotes Realta's opinion and is offered separately from the underlying trend analysis.

About This Research

Realta Research & Intelligence

Realta Research & Intelligence is our market research and thought leadership program committed to providing practical analysis to leaders responsible for fraud prevention, financial crime, payments risk, merchant onboarding, underwriting, compliance, and operational decisioning.

Our analysis combines publicly available industry research, regulatory developments, enforcement actions, and operational observations from across the payments ecosystem to help organizations understand where risk is changing—and what those changes mean in practice.

Realta Research & Intelligence publishes flagship market briefings, executive insights, and practical guidance. Rather than simply reporting industry developments, we focus on the operational implications for organizations making risk decisions across the payment lifecycle.

Topics include:

- Fraud and scam trends
- Financial crime and AML
- Payment risk
- Merchant onboarding and underwriting
- AI in fraud prevention
- Cyber-risk convergence
- Regulatory developments
- Connected decisioning

Our goal is straightforward:

Separate signal from noise

To get the latest insights from the Realta Research & Intelligence team, visit www.RealtaRisk.com.