Intelligence Report

IDENTITY FRAUD INTELLIGENCE

Trends & Insights

What network-scale data reveals
H2 2025–H1 2026

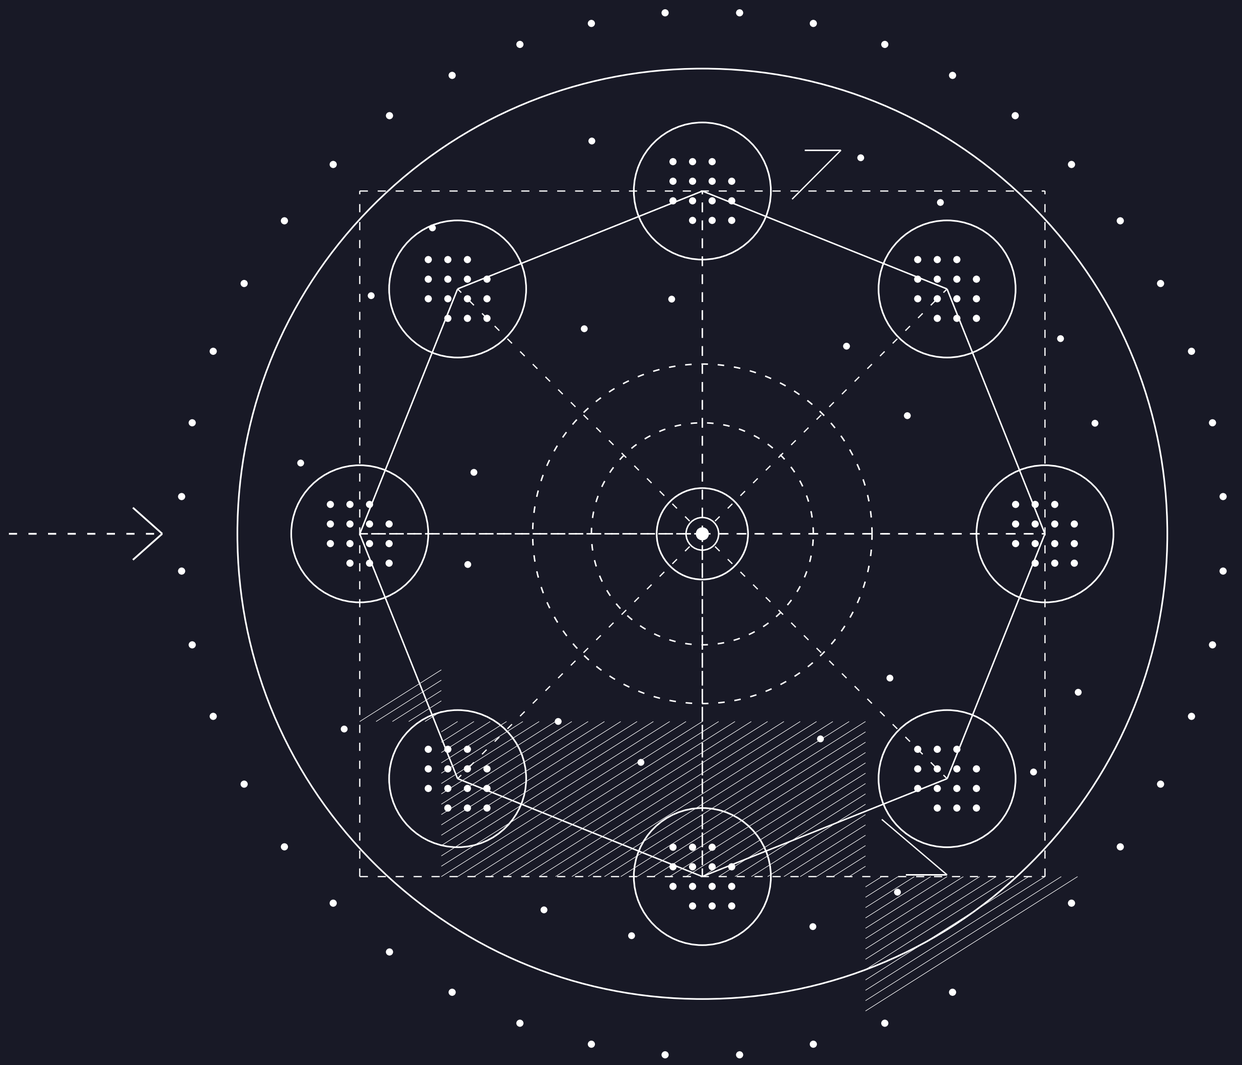

Liminal + UNICO

Contents

Navigation – click the Liminal + Unico logo to return to this page.

Executive Summary

Identity, long treated as a compliance checkbox, is now where revenue is protected or lost

01

Fraud has outgrown the cost of doing business

Global financial fraud losses now exceed $400 billion a year, and the force behind the number is generative AI. What used to be a predictable, budgeted cost is now growing faster than the systems meant to contain it.

02

The threat is not just growing. It is climbing.

Unico's network-scale data sorts attacks into three tiers. The crude tier is shrinking. Simple injection is the largest at 45.8 percent. The advanced tier already accounts for 23.3 percent, and the cost of reaching it has collapsed by more than two orders of magnitude.

03

There is no safe market, only different points on the same curve

Synthetic identity fraud, now a rising share of global cases, runs under every market. Latin America sits at the leading edge with the highest synthetic share in the world. It previews where other markets are heading; it is not a regional outlier.

04

The same actor already knows all of you

A single entity in Unico's network has been tied to 949 distinct identity documents, and one observed actor has targeted as many as 30 businesses. To an institution seeing only its own data, a repeat offender looks like a first-time visitor every time.

05

What stops being optional

Over the next eighteen months, sophisticated fraud becomes the leading tier, synthetic attacks become the default, and isolated defense becomes a structural liability. The market's own buyers are already moving, abandoning static checks for defenses that adapt in real time.

01
THE MAGNITUDE

Fraud has outgrown the cost of doing business

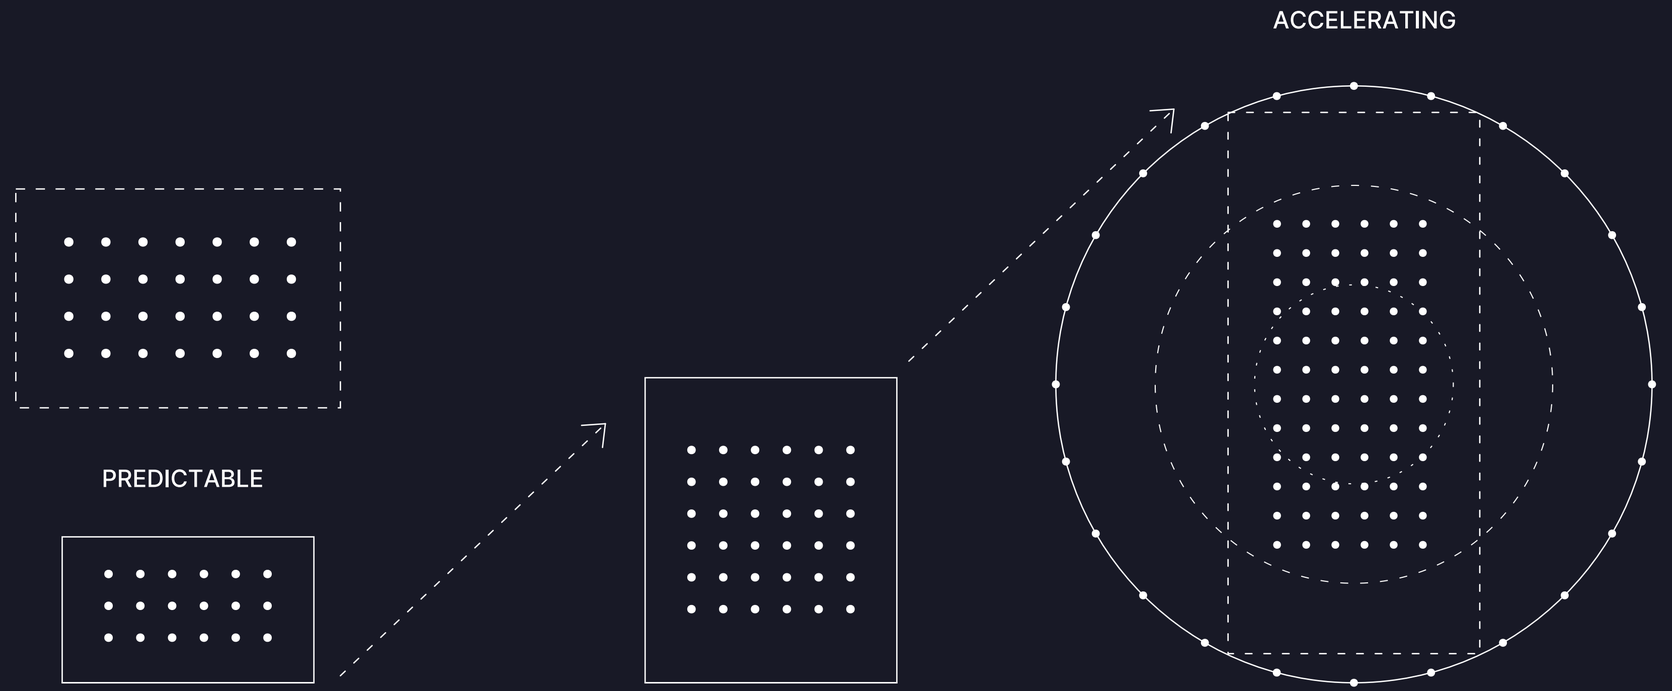

$347M

$400B+

121%

01
THE MAGNITUDE

Fraud has outgrown the cost of doing business

$347M

Chargeback budgets have grown 13% over the past two years

$400B+

Estimated global financial fraud losses in a single year, AI-fueled

121%↗

Projected rise in FI fraud losses by 2030

For years, fraud was a line item. Institutions modeled it, reserved against it, and treated a predictable slice of loss as the price of operating at scale. That logic held as long as fraud stayed expensive to commit and bounded in volume. It no longer does.

In early 2025, a single institution recorded 8,065 deepfake fraud attempts in eight months, with $347 million in verified losses traced to them.[1] No reserve quietly absorbs a number like that. It is a structural cost, not a line item a budget can plan around.

The aggregate is larger still. Industry estimates put global financial fraud losses past $400 billion in a single year, with generative AI driving the climb.[2] Treat that as the order of magnitude, not a precise count. The weight of this report rests on measured evidence, the verified losses above and the network-scale findings in the sections that follow, and the estimate only sets the scale around them. The mechanism, though, is not in dispute. The same tools that let a legitimate business scale its operations let an attacker scale a campaign at a marginal cost approaching zero. What once required money, equipment, and specialized skill now requires a subscription.

The people closest to the problem read it the same way. Nearly 90 percent of fraud practitioners now name AI-generated synthetic content as a leading concern, with manipulated documents and bot-driven impersonation the forms they find most alarming.[3]

The trajectory is not slowing. The same forecast projects financial-institution fraud losses to rise about 121 percent in five years, from roughly $25 billion in 2025 to more than $55.3 billion by 2030.[4] Markets do not move that way on their own. They move that way when the people attacking them are getting better, faster, and cheaper at once.

What changes here is the logic of the loss. When losses were predictable, absorbing them was rational, and a reserve worked as a buffer. When losses accelerate, compound, and increasingly arrive through attacks the institution never sees coming, that same reserve stops working as a buffer. It becomes a subsidy paid to whoever is on the other side.

02
THE SOPHISTICATION LADDER

The threat is not growing. It is climbing.

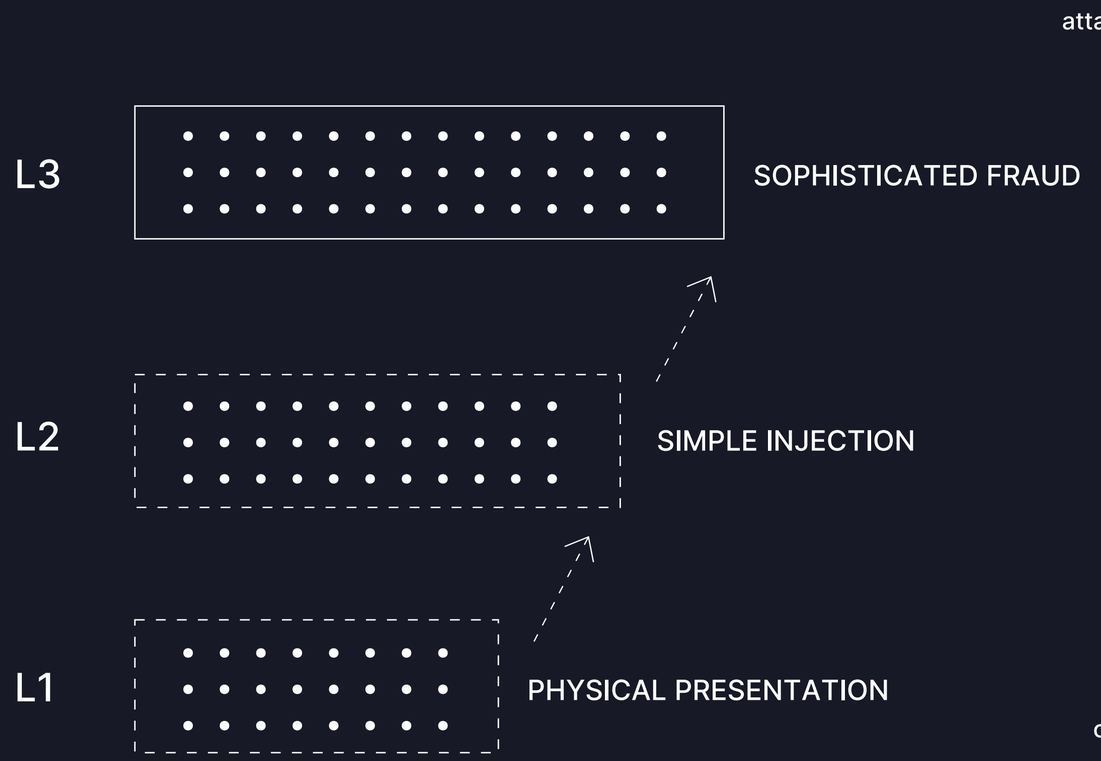

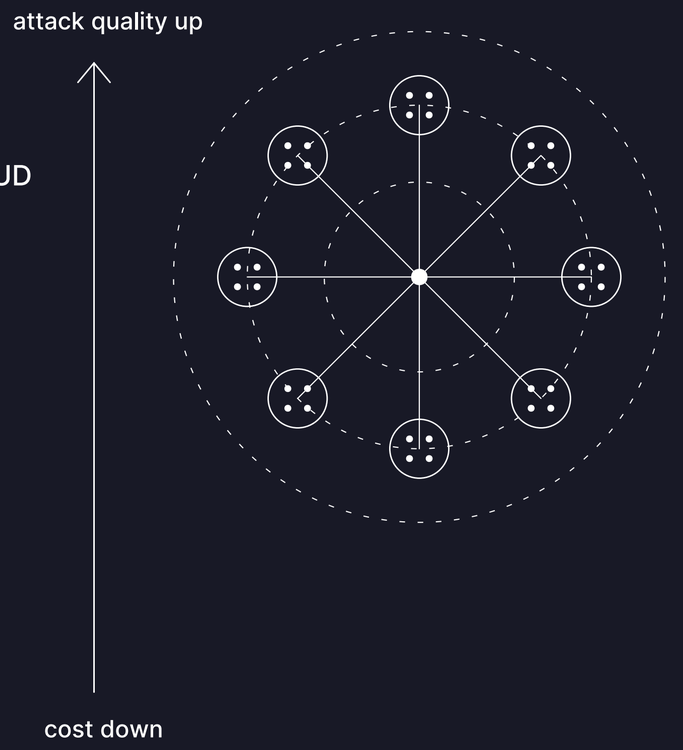

02
THE SOPHISTICATION LADDER

The threat is not growing. It is climbing.

Most fraud reporting counts volume. More attempts, more losses, more incidents year over year. Volume is the easy story, and it misses the harder one: the threat is getting better, not just bigger.

A count taken once is a snapshot of a single moment. Measuring the same network across nearly a year shows the direction and speed at which the threat moves, which a snapshot cannot. That vantage is rare, and it is what turns a static reading into a trend line.

Unico's Annual Fraud Sophistication Report, drawn from network-scale data gathered over roughly eleven months, June 2025 to April 2026, sorts classified attacks into three tiers by sophistication.

The fraud sophistication ladder

Tier	Attack type	Share of classified attacks
L3	Sophisticated fraud (deepfakes, injection, physical manipulations such as curved monitors and mirrors)	23.3%
L2	Simple injection (a manipulated image or video stream fed directly into the verification pipeline)	45.8%
L1	Physical presentation (a printed photo or a face held to a camera)	30.9%

Shares of classified attacks, June 2025 to April 2026. Source: Unico Annual Fraud Sophistication Report.

02

THE SOPHISTICATION LADDER

The threat is not growing. It is climbing.

~9x

Growth in monthly classified injection cases across the window (Unico)

>100x

Drop in the cost of executing a sophisticated attack versus a few years ago (Unico)

Physical presentation, the crude tier, a printed photo or a face held to a camera, still accounts for 30.9 percent: still material, but declining. For any institution running modern liveness detection, these crude attacks are increasingly contained, and the share keeps shrinking.

The middle tier is where the volume sits. Simple injection, where an attacker feeds a manipulated image or video stream directly into the verification pipeline, is the single largest category at 45.8 percent. It is best read as a waiting room: the tier sitting directly below the top, filling fast, one cheap step from graduating upward.

The top tier is the one to watch. Sophisticated fraud, combining deepfakes, injection, and physical manipulations such as curved monitors and mirrors, already makes up 23.3 percent of classified attacks. A year ago that tier was a frontier. Today it is a quarter of everything.

The distribution is only part of the finding. The other part is motion, the piece a snapshot would miss. Two of Unico's numbers measure speed rather than level, and they are what make the climb undeniable.

What moves an attacker up the ladder is cost, and that cost has collapsed by more than two orders of magnitude. The barrier that once kept the top tier small has effectively been removed, and every actor in the 45.8 percent middle is one cheap step from the 23.3 percent above it.

Liminal Intelligence places this climb on a longer curve. Fraud losses at financial institutions are forecast to rise about 121 percent by 2030, and the composition of those losses shifts toward the advanced tiers as the cost of sophistication keeps falling.[4] The ladder describes a direction of travel, not a fixed snapshot, and every force acting on it points up.

The harder challenge is structural. In Liminal's research, 93 percent of practitioners express confidence in their fraud models, while 92 percent acknowledge that legacy infrastructure still lets fraud signals slip through, with manual review, rule-based systems, and even AI and machine-learning tools all showing frequent misses.[5] The gap opens precisely where defenses feel most settled, because the tooling was built for a threat that has since moved.

Set the two forces side by side and the economic case is plain: the cost of mounting a sophisticated attack has fallen more than 100x (Unico) while the cost of defending in isolation is projected to rise about 121 percent by 2030 (Liminal), and the point where the second overtakes the first is the point where isolation stops paying for itself.

03
THE GLOBAL CURVE

There is no safe market, only different points on the same curve

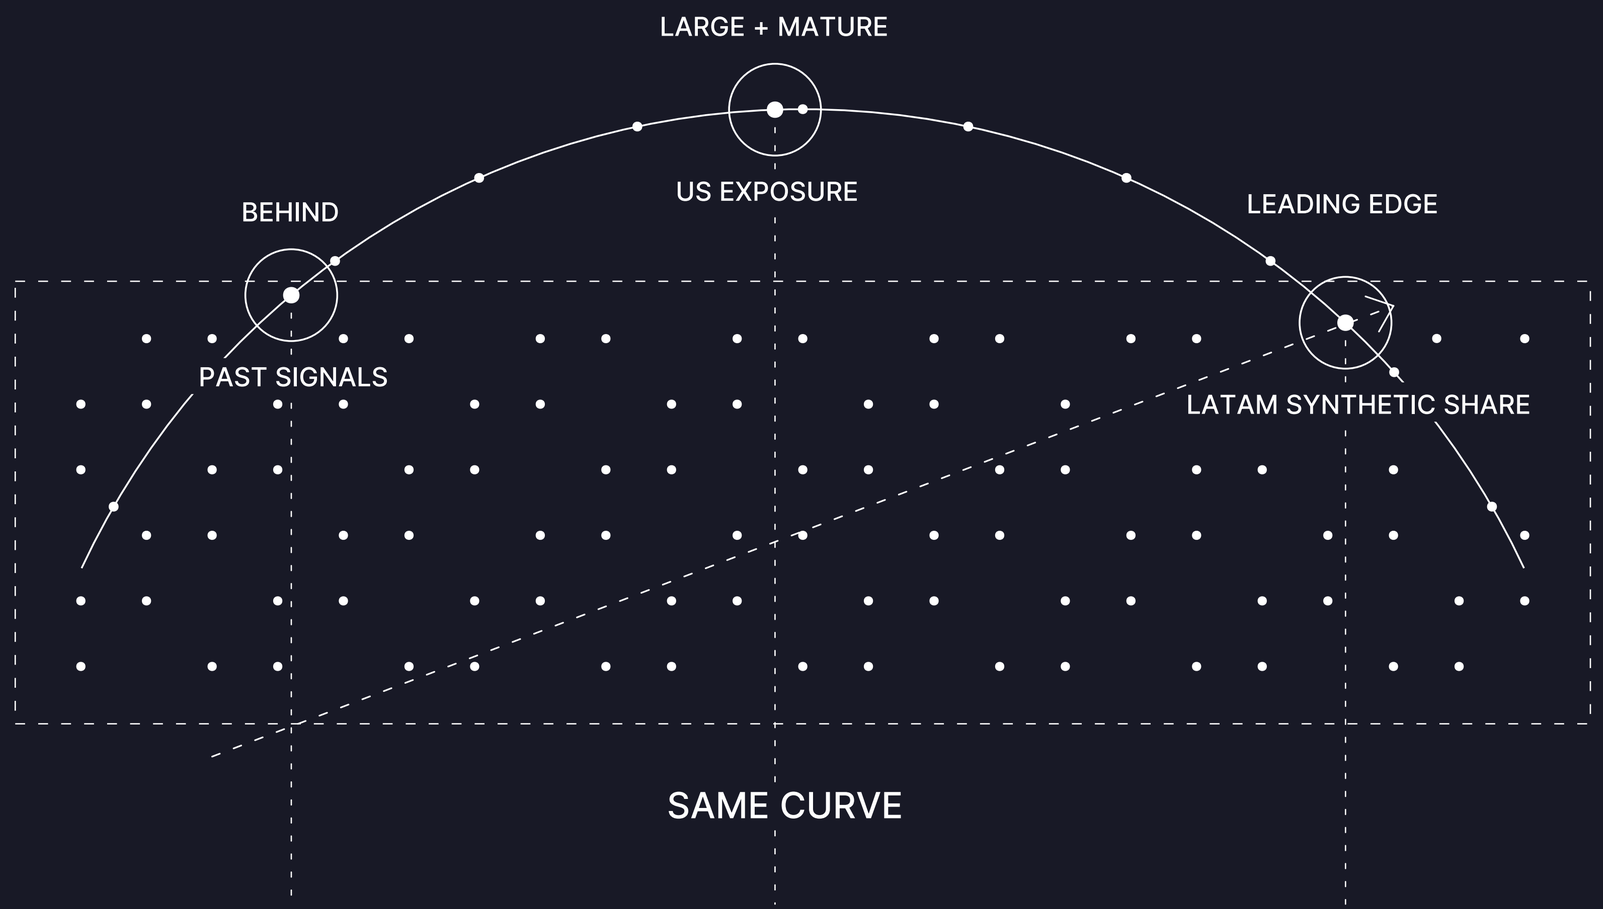

03

THE GLOBAL CURVE

There is no safe market, only different points on the same curve

48.3%

LATAM synthetic fraud share, highest in the world

6M

Account-takeover victims in 2025, up 141% since 2021

$280K+

Average loss per successful deepfake-enabled incident

It is tempting to read fraud as a local problem. Every market has its own regulators, its own payment rails, its own attackers, and the instinct is to treat the threat as something shaped by geography. The data complicates that instinct. The techniques are global, the tooling is borderless, and the same curve runs underneath every market. What differs between markets is not the curve but the position on it.

Synthetic identity fraud is the clearest case. It now accounts for 11 percent of all global fraud cases, a figure that has grown roughly eightfold year over year.[6] This is the fraud with no victim to raise an alarm, because the identity it uses was never real. It is assembled rather than stolen, and it is the category climbing fastest worldwide.

The pattern is not confined to one technique or one region. Global fraud attacks rose 8 percent in the most recent measure, account takeover affected roughly six million victims in 2025, up 141 percent since 2021, and the average successful deepfake-enabled incident now costs an organization more than $280,000, with nearly one in five reporting losses above half a million.[7] That figure counts only attacks that land: it is not comparable to the $347 million cited earlier, which averages across 8,065 attempts, most of them blocked. These figures describe the global baseline, and every market sits somewhere on it.

The threat is broad-based rather than concentrated in a single method. In Liminal's demand-side research, card misuse (21 percent), card testing (17 percent), account takeover (17 percent), and triangulation schemes (14 percent) each carry a comparable slice, spreading strain across every defense at once.[8]

Latin America is the leading edge of that baseline. Synthetic identity fraud makes up 48.3 percent of the region's fraud cases, the highest share recorded anywhere in the world.[9] It would be comfortable to explain that number away as a regional anomaly. It is closer to a preview. LATAM reached this point first because it combined fast digital adoption, high mobile penetration, and rapid financial inclusion, the same conditions now spreading across other emerging markets. The volume is already arriving at scale: in a single recent deployment, more than 500,000 deepfake identity attempts were blocked across banks and fintechs in Brazil and the wider region.[10] Where LATAM is today, other regions are heading.

A market further along in its defenses has not left the curve; it occupies a different point on it. In the United States, synthetic identity fraud drove an estimated $3.3 billion in exposure at financial institutions in 2024, even with deeper defenses and tighter regulation in place.[11] The share is lower than LATAM's, but the absolute exposure is enormous. Maturity moves a market further along the curve; it does not exempt the market from it.

Exposure to the techniques is near-universal. What sets one market's position apart from another's is a small set of structural variables: the speed of digital adoption, mobile penetration, the pace of financial inclusion, regulatory maturity, and how widely modern liveness and connected verification are already deployed. Read those variables together and a market's place on the curve becomes legible, including the climb it has not yet seen.

03
THE GLOBAL CURVE

There is no safe market, only different points on the same curve

None of this means the threat is evenly distributed; it plainly is not. The useful question for any single institution is how far up the curve its market already sits, and how ready it is for the stretch it has not yet climbed.

This is where comparison becomes the most useful instrument an institution has. An organization that benchmarks itself only against its domestic peers is measuring against others at the same point on the curve. The markets ahead of it have already lived through the threats it is about to meet. The ones behind it show what its own recent past looked like. Read together, they map the threats a market has not yet met.

→ WHAT THIS MEANS FOR YOU

The right benchmark is not your domestic peer set. It is the market one step ahead of you on the curve, because that is the threat arriving at your door next.

Positions on the global fraud curve

Position	What defines it	Where the data places it
Leading edge	High-velocity emerging markets: fast digital adoption, high mobile penetration, and rapid financial inclusion running ahead of defense maturity	LATAM: synthetic identity fraud at 48.3% of cases, the highest share in the world; 500,000+ deepfake attempts blocked across Brazil and the region in a single recent push
Large and mature	Deep defenses and tighter regulation, but vast transaction volume and legacy identity rails	United States: an estimated $3.3 billion in synthetic identity exposure at financial institutions in 2024, large in absolute terms even where the share is lower
What sets the position	The variables that place any market on the curve	Speed of digital adoption, mobile penetration, pace of financial inclusion, regulatory maturity, and the reach of modern liveness and connected verification

Selected positions, not an exhaustive ranking. Sources: LexisNexis Risk Solutions; DuckDuckGoose; TransUnion; Liminal analysis.

Liminal + UNICO

04
THE NETWORK

The same actor already knows all of you

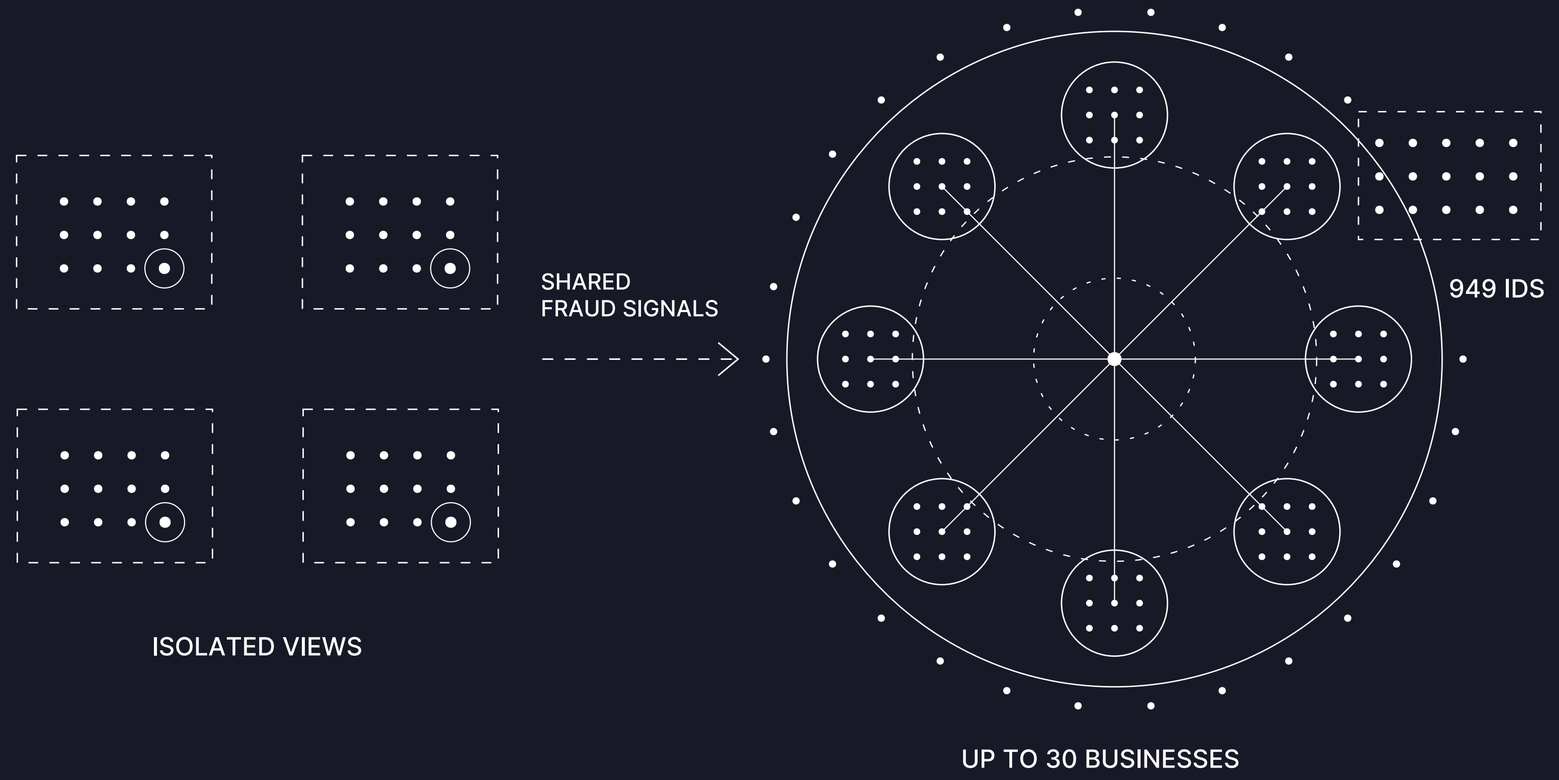

04

THE NETWORK

The same actor already knows all of you

So far the threat has been described from the outside: how big it is, how fast it climbs, how it spreads across markets. Change the vantage point now, and look at the whole network at once rather than one transaction at a time. That shift produces the most important finding in this report.

Fraud is easy to picture as isolated incidents. One fake identity, one attempt, one company. Unico's network data describes something else. A single evaluation entity in Unico's global base has been tied to 949 distinct identity documents, and that entity is not a one-off: the top five entities each carry more than 850.[12] At that scale it is not occasional fraud but an operation run like a business.

A single operator, rejected by one institution, simply moves to the next, and the next, carrying the same identities into each one. Nothing about the rejection follows them.

The same pattern holds across companies, not only identities. Unico has observed one actor target as many as 30 different businesses.[13] Each institution meets that actor fresh. This is the stranger problem: an isolated institution greets every repeat offender as a first-time visitor.

It is not only Unico's vantage. The institutions on the other side report the same shape from below, with nearly all practitioners, 92 percent, saying third-party fraud spans multiple accounts or platforms.[14]

The institutions on the receiving end feel this as repetition. In Liminal's research, 85 percent of organizations report repeat account takeover incidents, nearly half of them frequently, because once an account is compromised, attackers return using the same techniques.[15] That is the same actor the network data identifies from above: one side counting 949 identities behind a single entity, the other counting the same offender arriving again and again.

The structural flaw is simple to state. When every institution sees only its own slice, a repeat offender looks like a first-time visitor each time they cross from one organization to the next. The 949th attempt and the first attempt look identical to a company working from its own data alone. The actor blocked by company one arrives at company two already practiced, already funded, and completely unknown. Isolation does not make an institution a smaller target. It makes it a blind one.

The same logic runs in reverse. When verification is connected, the fraudster's history travels with them, and the actor rejected by company one arrives at company two already flagged. The 949th attempt stops resembling the first, because the network remembers what a single institution cannot. Fraud operates as a network. A defense that does not operate as one is fighting an organized adversary one isolated case at a time.

The property that makes a connected network powerful makes it hazardous if built carelessly. If a fraudster's history travels, so does a mistake. A legitimate customer wrongly flagged at one institution can arrive at the next already marked, rejected across many businesses for a judgment that was wrong the first time. That is the 949-identity finding turned against the innocent, and it is the risk an executive should weigh before connecting anything. It also defines what separates a credible network from a reckless one: not the signals it shares, but the correction path it builds, a way for a wrongful flag to be challenged, traced, and cleared across the network rather than only recorded across it. Isolated defense offers no recourse either. A customer wrongly rejected today simply vanishes, and the institution never learns it erred. Connected defense does not eliminate the false flag. It is the only model that can see one at all, and so the only one that can correct it. So the choice is narrower than it looks: share signals only alongside a way to catch a wrongful flag and reverse it across the network.

04
THE NETWORK

The same actor already knows all of you

Buyers are already funding the shift. Strengthening partnerships now ranks among the top investment priorities for the year ahead, at 47 percent, and shows the largest year-over-year increase of any category.[16] The demand side is moving its budget toward exactly the connected posture the network data argues for.

The market has started to act. Over the past year, cross-border fraud detection has moved from concept to live infrastructure: shared intelligence networks now link institutions across regions, a major cross-border payments body has begun pooling fraud signals across its members, and live cross-institution networks have launched in multiple markets.[17] The organizations furthest ahead have already concluded that no institution can see enough on its own.

→ WHERE DOES YOUR INSTITUTION SIT?

Find the rung that describes how your verification works today, not how you wish it worked.

Rung	Stage	How verification works
01	Isolated and static	Identity is checked once, at onboarding, against a document. Defense is point-in-time and rules-based. The institution sees only its own data. Every repeat offender arrives as a stranger.
02	Reactive	Liveness and fraud signals are added after attacks succeed. Detection is still mostly retrospective, reviewed after the fact. Crude attacks are caught; injection and synthetic identities mostly are not.
03	Continuous, internal	Strong biometric verification and liveness detection anchor the defense, catching the presentation, injection, and deepfake attacks the sophistication ladder is built on, with behavioral and device signals adapting in real time. This is the layer that stops a sophisticated attack at the door. Its limit is reach, not strength: visibility stops at the institution's own perimeter, so the stranger problem persists. You see your own repeat offenders, never the network's.
04	Connected	Strong biometric and liveness defense now works in tandem with a collective network. Biometrics confirm the person at the door; the network confirms whether that identity has been seen, flagged, or rejected elsewhere. A fraudster's history begins to travel, and the 949th attempt no longer resembles the first. Neither layer replaces the other: a network without biometric defense is blind to the attack in front of it, and biometric defense without the network is blind to every other door the actor has used.
05	Network-native	Biometric defense and the network operate as one shared, adaptive layer. An actor blocked at one institution arrives at the next already known, and the same liveness checks that catch a novel attack feed what the network learns. Defense moves at the speed of the attacker because it sees what the attacker sees: the whole field.

Most institutions today sit at Rung 2 or 3. The threat described in this report operates at Rung 5. **The gap between where you sit and where the threat operates is the single best predictor of your exposure over the next eighteen months.**

05
THE FORWARD OUTLOOK

What stops being optional

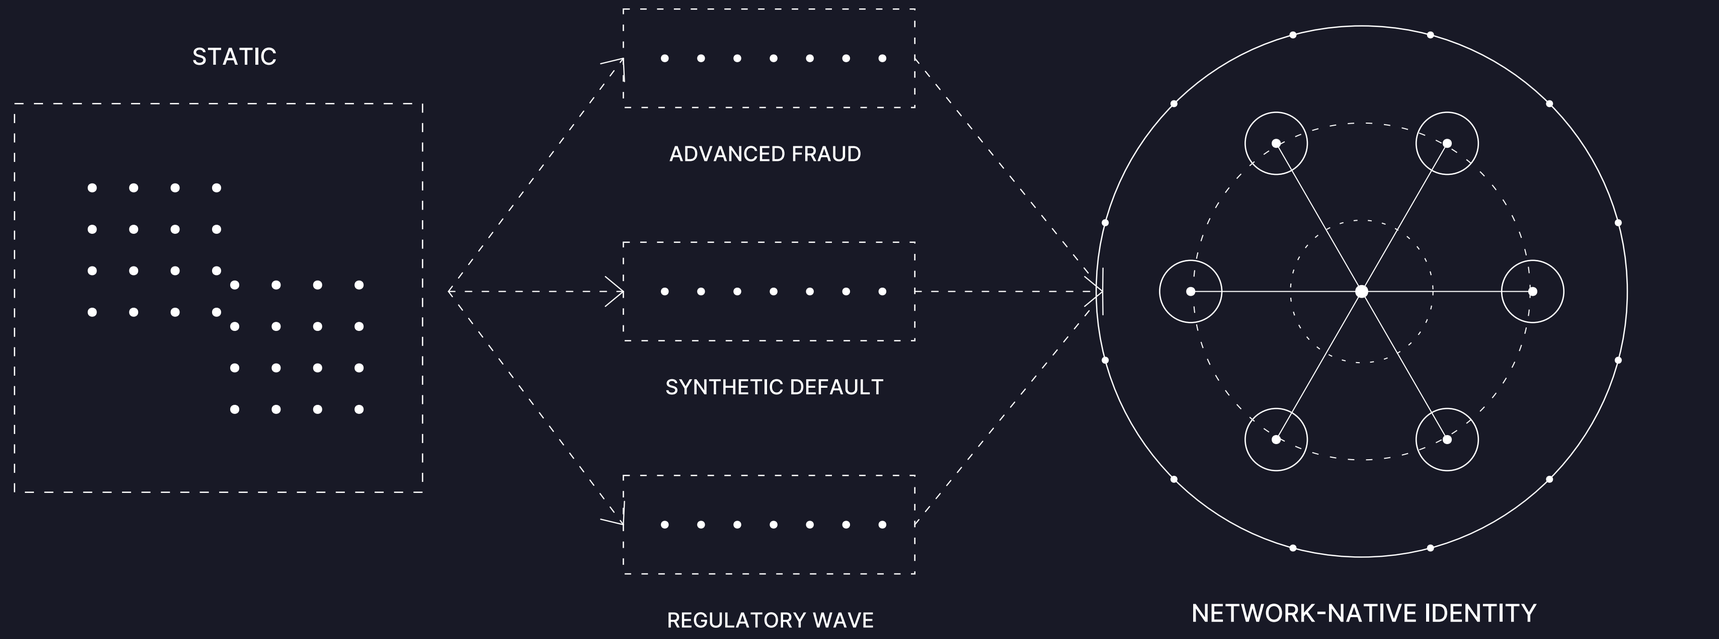

05

THE FORWARD OUTLOOK

What stops being optional

A forecast is only as good as the mechanism under it. What follows is not speculation about where fraud might drift; each outlook is the visible continuation of a force already measured in the sections above. The open question is how much of the next eighteen months an institution spends ready for these shifts rather than surprised by them.

The sophistication ladder inverts

Sophisticated fraud, today roughly a quarter of classified attacks, is positioned to become the leading tier within twelve to eighteen months. The mechanism is already established: the cost of advanced tooling has collapsed by more than two orders of magnitude, and the largest tier today, simple injection at 45.8 percent, sits one short skill-step below the top.[18] The population able to graduate upward is already large and already active. Liminal Intelligence places the same trajectory on the loss curve, with financial-institution fraud losses forecast to rise about 121 percent by 2030, from roughly $25 billion to more than $55.3 billion.[4] What is the frontier today becomes the baseline within the forecast window.

Crude fraud fades to residual

The bottom rung continues to erode. Physical-presentation attacks, photos and screens held to a camera, keep shrinking toward a small residual share for any institution running modern liveness.[18] This is the one piece of unambiguously good news in the forecast, and it carries a warning inside it. The floor is shrinking because the cheap attacks are migrating up the ladder, not because fraud is in retreat.

Synthetic and AI-enabled fraud becomes the dominant mode of attack

The technique with no real victim to raise an alarm is already the fastest-growing category worldwide. At 11 percent of global cases today, growing roughly eightfold year over year, and already 48.3 percent in the leading market, the trajectory points toward dominance well within the forecast window. In Liminal's research, practitioners project synthetic identity fraud (79 percent) and mule networks (69 percent) to accelerate faster than any other vector, 74 percent name detection of synthetic or first-party activity a significant weakness, and 88 percent have already seen synthetic content feeding third-party schemes.[19] The conditions driving it, cheap generative tooling and borderless distribution, are not market-specific. What is concentrated in the leading markets today becomes the global baseline over the forecast window.

Isolated defense becomes untenable

This is the outlook that ties the others together. As attacks grow more sophisticated, more synthetic, and more networked, the single-institution view becomes a structural liability rather than a manageable constraint. The market has already begun to move. In Liminal's research, buyers are shifting decisively away from static, point-in-time checks toward continuous, behavioral, and device-based defenses that adapt in real time; 76 percent now rank device fingerprinting as their most effective defense against AI-enabled fraud, while document verification and other one-time controls fall to the back.[20] The same buyers are putting budget behind it: 62 percent expect to raise compliance spend in the year ahead.[16] Collective, connected defense is following the same path, with cross-border and cross-institution fraud networks moving from concept to live infrastructure inside the last year.[17] The institutions that treat this as optional are deciding to meet an organized adversary alone.

05
THE FORWARD OUTLOOK

What stops being optional

Regulation turns adaptive, auditable identity from a choice into an obligation

The economic case is one source of pressure. The compliance case is another, and it arrives on a fixed timeline. The EU AI Act, the only horizontal AI law from a major economy, entered into force in 2024 and binds in stages: the first prohibitions landed in early 2025, and the high-risk obligations covering biometric identification take effect in December 2027.[21] It does not stand alone. Across 2027 a cluster of EU regimes becomes binding on banks, payment providers, and identity vendors at once: the single AML rulebook, the digital identity wallet framework, the payments package with stronger customer authentication, and the operational resilience regime already in force. Liminal reads them as converging on one endpoint, credential-presented and auditable identity as the gating layer for participation in regulated digital economies.

The fraud-detection carve-out that vendors point to is, in Liminal's assessment, narrower than it is sold to be: the moment a fraud signal drives a credit denial or an account restriction, its exemption is contested.[22] As with the fraud curve, the EU is the leading regulatory edge rather than the whole story. Outside it the picture is more fragmented. The United States has no horizontal AI or privacy law, leaving a patchwork of state-level biometric and privacy rules. Brazil, where much of this report's network data originates, governs identity and data through the LGPD and an active open-finance agenda, without an equivalent binding AI timeline yet. The EU matters globally less because every market will copy it than because its rules reach any institution serving EU customers and tend to set the template others adapt. A defense that cannot explain, log, and adapt its identity decisions will be both weaker against the threat and harder to defend to a regulator.

The regulatory wave: what binds when

When	What binds
August 2024	EU AI Act enters into force
Early 2025	First AI Act prohibitions apply; the operational resilience regime (DORA) in force from January 2025
Across 2027	A cluster becomes binding together: the single AML rulebook, eIDAS 2.0 digital identity wallet acceptance, and the PSD3 and PSR payments package with stronger customer authentication
December 2027	AI Act high-risk obligations apply, including biometric identification

Selected EU milestones, some subject to final texts. Source: Liminal regulatory analysis of EU instruments. The EU is the leading regulatory edge; its rules reach institutions serving EU customers.

05
THE FORWARD OUTLOOK

What stops being optional

→ THE OBVIOUS OBJECTION: WE CANNOT SIMPLY SHARE OUR DATA

Every argument for connected defense runs into the same wall, and it is a real one. Institutions cannot freely pool customer data. Privacy regulation restricts it, competitive instinct resists it, and in many markets the law forbids it outright. An executive reading this is right to be skeptical of any conclusion that depends on banks opening their books to one another. That is not how the industry works, and it is not how it should work.

The objection is correct about the constraint and wrong about the conclusion. Connected defense does not require sharing customer data. The collective-defense infrastructure now moving from concept to live deployment is built precisely to avoid it: institutions exchange fraud signals, not customer records. A network can confirm that an identity has been seen, flagged, or rejected elsewhere without any institution surrendering its data, its customers, or its competitive position. What travels is the fact of the fraud, not the identity of the customer.

This is the distinction that makes the forward outlook viable rather than aspirational. The institutions furthest up the maturity ladder have not abandoned privacy or competitive boundaries. They have found the narrow path between them: shared intelligence about bad actors, held separately from protected data about good ones. The question for everyone else is not whether to surrender data. It is whether to keep meeting known offenders as strangers while a way to recognize them already exists.

Set the outlooks side by side and the next eighteen months take a single shape. Fraud is becoming more advanced, more synthetic, more global, and more networked at once, for one underlying reason: the economics now favor the attacker at every tier. The cost of mounting an advanced, synthetic, cross-institution attack has fallen to near zero, while the cost of defending against it alone has quietly climbed to the point of futility.

This is why the framing has to change. Identity verification has long been treated as a compliance function, a checkbox cleared before the real business begins. It now sits at the center of the business: the place where trust is won or lost, and where the most organized adversaries concentrate their effort. It has become a determinant of the business rather than a gate in front of it.

None of this prescribes a single answer. The right response will differ by market, by institution, and by where each one sits on the curve traced here. What the evidence settles is the direction. The threat is connected and adaptive, and the market's own buyers are moving the same way, shifting spend toward defenses that keep working past the first check. A static, isolated defense is built for a problem that no longer exists.

Sources & Notes

Sources & Notes

1. Deepfake fraud attempts (8,065) and $347M verified losses at a single institution, January to August 2025. As reported by Group-IB.
2. Generative-AI-enabled fraud losses exceeding $400 billion within a single year. As reported by TechRadar, March 2026, citing industry fraud analysis.
3. AI-generated content as a leading fraud concern: nearly 90% of practitioners cite synthetic content as a leading concern, with manipulated documentation (72%) and bot-driven impersonation (67%) most alarming. Liminal, Link Index for Consumer Payment Fraud in Banking, 2026 (Third Party Fraud Buyer Demand Survey, August 2025, N=72).
4. Financial-institution fraud losses forecast to rise from less than $25 billion (2025) to more than $55.3 billion by 2030. Global fraud and risk outlook, April 2026, via Liminal market intelligence.
5. Confidence and legacy gap: 93% of practitioners express confidence in their fraud models, while 92% acknowledge missing fraud signals due to legacy infrastructure, with manual reviews (58%), rule-based systems (57%), and AI and machine-learning tools (55%) all showing frequent misses. Liminal, Link Index for Consumer Payment Fraud in Banking, 2026 (Third Party Fraud Buyer Demand Survey, August 2025, N=72).
6. Synthetic identity fraud as a share of global fraud cases and year-over-year growth. LexisNexis Risk Solutions, April 2026.
7. Global fraud-attack rise (8%); account-takeover victims (6M, up 141% since 2021); average deepfake incident loss ($280,000+, with detail on losses above $100,000 and $500,000). Attack-rise and account-takeover figures: LexisNexis Risk Solutions, 2025 to 2026, via Liminal market intelligence. Per-incident deepfake loss: IRONSCALES, Deepfake Impact Survey, November 2025.
8. Fraud-type distribution: card misuse (21%), card testing (17%), account takeover (17%), and triangulation schemes (14%). Liminal, Link Index for Consumer Payment Fraud in Banking, 2026 (Third Party Fraud Buyer Demand Survey, August 2025, N=72).
9. Latin America synthetic identity fraud share (48.3%), highest recorded globally. LexisNexis Risk Solutions, April 2026.
10. More than 500,000 deepfake identity attempts blocked across banks and fintechs in Brazil and Latin America in a single recent deployment. DuckDuckGoose, February 2026.
11. Synthetic identity fraud exposure at US financial institutions of an estimated $3.3 billion in 2024. TransUnion.
12. Single evaluation entity associated with 949 distinct identity documents; top five entities each exceeding 850. Unico, Annual Fraud Sophistication Report (June 2025 to April 2026).
13. Single observed fraud actor targeting as many as 30 distinct businesses. Unico, Annual Fraud Sophistication Report (June 2025 to April 2026).
14. Nearly all practitioners (92%) report third-party fraud spanning multiple accounts or platforms. Liminal, Link Index for Consumer Payment Fraud in Banking, 2026 (Third Party Fraud Buyer Demand Survey, August 2025, N=72).
15. 85% of organizations report experiencing repeat account-takeover incidents at least occasionally, with 48% occurring frequently. Liminal Buyer Survey, March 2025, published in the 2025 Financial Crime & Compliance Market & Investment Trends Report (March 2026).
16. Investment priorities and budgets: strengthening partnerships is among the top 2026 investment priorities at 47% and shows the largest year-over-year increase of any category; 62% of organizations expect compliance budget increases in 2026. Liminal Buyer Survey, March 2025, published in the 2025 Financial Crime & Compliance Market & Investment Trends Report.

Sources & Notes

17. Collective-defense market signals: cross-border and cross-institution fraud-detection networks moving to live infrastructure over the prior year. Liminal market intelligence, compiled from public announcements, 2025 to 2026.

18. Fraud-tier distribution, ninefold growth in monthly injection cases, and the more than 100x collapse in the cost of executing a sophisticated attack. Unico, Annual Fraud Sophistication Report (June 2025 to April 2026).

19. Synthetic projections and detection gaps: synthetic identity fraud (79%) and mule networks (69%) projected to accelerate fastest; 74% identify detection of synthetic or first-party activity as a significant weakness; 88% have observed synthetic content contributing to third-party schemes. Liminal, Link Index for Consumer Payment Fraud in Banking, 2026 (Third Party Fraud Buyer Demand Survey, August 2025, N=72).

20. Buyers shifting from static, point-in-time checks to continuous, behavioral, and device-based defenses; 76% rank device fingerprinting as the most effective defense against AI-enabled fraud, ahead of behavioral biometrics (64%) and device intelligence (61%), with document verification lagging. Liminal AML Market & Buyers Guide, June 2025, published in the 2025 Financial Crime & Compliance Market & Investment Trends Report.

21. EU regulatory timeline: the EU AI Act (Regulation (EU) 2024/1689) entered into force August 2024, with first prohibitions from February 2025 and high-risk obligations covering biometric identification binding December 2027; alongside the single AML rulebook (AMLR), eIDAS 2.0 digital identity wallet acceptance, the PSD3 and PSR payments package with stronger customer authentication, and the DORA operational resilience regime (in force since January 2025). Liminal regulatory analysis, 2026, drawing on EU primary sources.

22. Liminal analytical read: the AI Act fraud-detection carve-out (Annex III(5)(b)) is narrowly drawn and contested where a fraud signal drives credit denial or account restriction; the Commission’s draft Article 6(3) guidelines (May 2026) signal a narrow interpretation, with profiling explicitly excluded. Liminal analysis.

Bibliography

Unico. Annual Fraud Sophistication Report, June 2025 to April 2026. Primary source for the fraud sophistication ladder, injection growth, cost collapse, and the network repeat-offender findings (949 identities; up to 30 businesses).

Liminal. 2025 Financial Crime & Compliance Market & Investment Trends Report, March 2026. Source for the buyer shift toward continuous, behavioral, device-based defense (76% device fingerprinting) and repeat account-takeover incidence (85%, with 48% frequent). Underlying data: Liminal Buyer Survey, March 2025; Liminal AML Market & Buyers Guide, June 2025.

Liminal. Link Index for Consumer Payment Fraud in Banking, 2026. Source for synthetic identity and mule-network acceleration and practitioner-identified detection gaps (74%).

IRONSCALES. Deepfake Impact Survey, November 2025. Source for average per-incident deepfake loss exceeding $280,000 and the distribution of losses above $100,000 and $500,000.

DuckDuckGoose. Deepfake detection deployment, February 2026. Source for deepfake identity attempts blocked across banks and fintechs in Brazil and Latin America (500,000+).

TransUnion. Synthetic identity fraud analysis, 2024. Source for US financial-institution synthetic identity exposure (an estimated $3.3 billion, 2024).

Liminal. EU regulatory research, including The EU Regulatory Wave: What Binds When and The AI Act for Banks: Deployer Obligations, 2026. Source for the EU regulatory timeline and the convergence of the AI Act, the AML rulebook, eIDAS 2.0, PSD3 and PSR, and DORA, and for the analytical read on the AI Act fraud-detection carve-out. Draws on EU primary sources, including Regulation (EU) 2024/1689.

LexisNexis Risk Solutions. Global fraud trends, April 2026. Source for synthetic identity fraud as a share of global cases, year-over-year growth, and the Latin America regional share (48.3%).

Global fraud and risk outlook (industry analysis), April 2026. Source for the financial-institution loss forecast ($55.3 billion by 2030).

Group-IB; TechRadar coverage, 2025 to 2026. Sources for single-institution deepfake-loss figures and the $400 billion single-year global estimate.

Methodology Note

Figures attributed to secondary coverage were cross-checked against Liminal's market intelligence corpus where possible. Where an originating research firm could not be independently confirmed, the reporting outlet is cited. Unico network figures derive from Unico's own evaluation base over the June 2025 to April 2026 observation window.

The Actionable Intelligence Company

The same depth and rigor you've just read, now powering the verified intelligence behind your next strategic decision

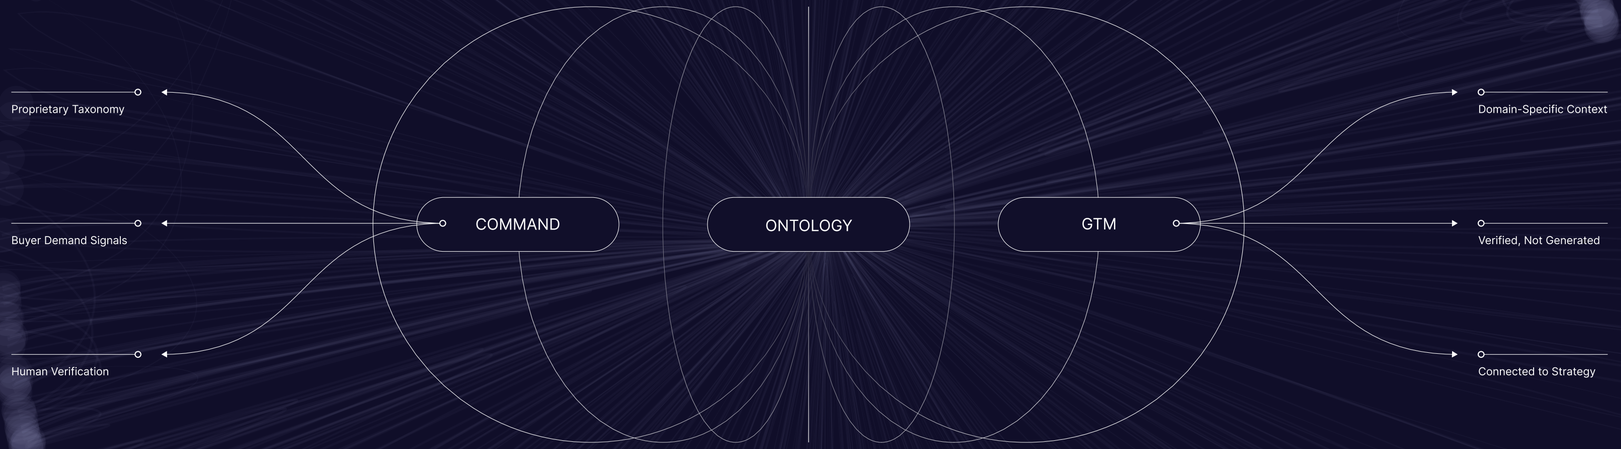

Three interfaces, one verified foundation:

Command

validate roadmaps against verified buyer demand, monitor competitive moves, and detect market shifts before they go public

GTM

target the right accounts, arm reps with live battlecards and AI-coached outreach, and close with context no competitor can access

Expert Desk

every signal verified by human domain experts recruited from Identity, Fraud, and Cybersecurity

Ontology

All three run on a continuously verified ontology of 2.5M+ real relationships across vendors, buyers, regulations, capabilities, and threat vectors.

Liminal maps the real relationships in Identity, Fraud, and Cybersecurity into a continuously verified, queryable ontology. Purpose-built interfaces sit on top: Command for strategy and product teams, GTM for revenue teams, and the Expert Desk for human verification.

Proven
across our portfolio

9x
FASTER STRATEGIC DECISIONS

30%
HIGHER COMPETITIVE WIN RATES

70%
LESS MANUAL RESEARCH TIME

From strategy to execution, verified. See how one intelligence system powers both the decisions you make and the deals you win.

GET ACCESS

unico

The truth layer for the GenAI era

Unico is an identity verification and fraud prevention platform that delivers deterministic, binary answers — the user is who they claim to be. No scores to interpret, no thresholds to configure, no manual review queues.

Built on 17 years of proprietary data and 3B+ verifications processed, Unico's network-effect model makes every verification smarter than the last. The platform combines real-time biometric verification, deepfake and spoofing detection, fraud risk classification — all in under 3 seconds.

Unico protects critical moments across the customer lifecycle: onboarding, account recovery, high-value transactions, payment authorization, and multi-account detection. One selfie. One process. One result.

Explore Unico's full identity and fraud prevention platform **unico.io**

Unico by the numbers

$7B+
in financial fraud prevented across the platform

400+
companies across 25+ countries

150M+
transactions per month processed

99.9%
approval rate — 97% on the first attempt

Intelligence Report

IDENTITY FRAUD INTELLIGENCE

Trends & Insights

What network-scale data reveals
H2 2025–H1 2026

WWW.LIMINAL.CO WWW.UNICO.IO INFO@LIMINAL.CO

Liminal + Unico