PLAID

The new identity crisis:

Rethinking fraud in the AI era

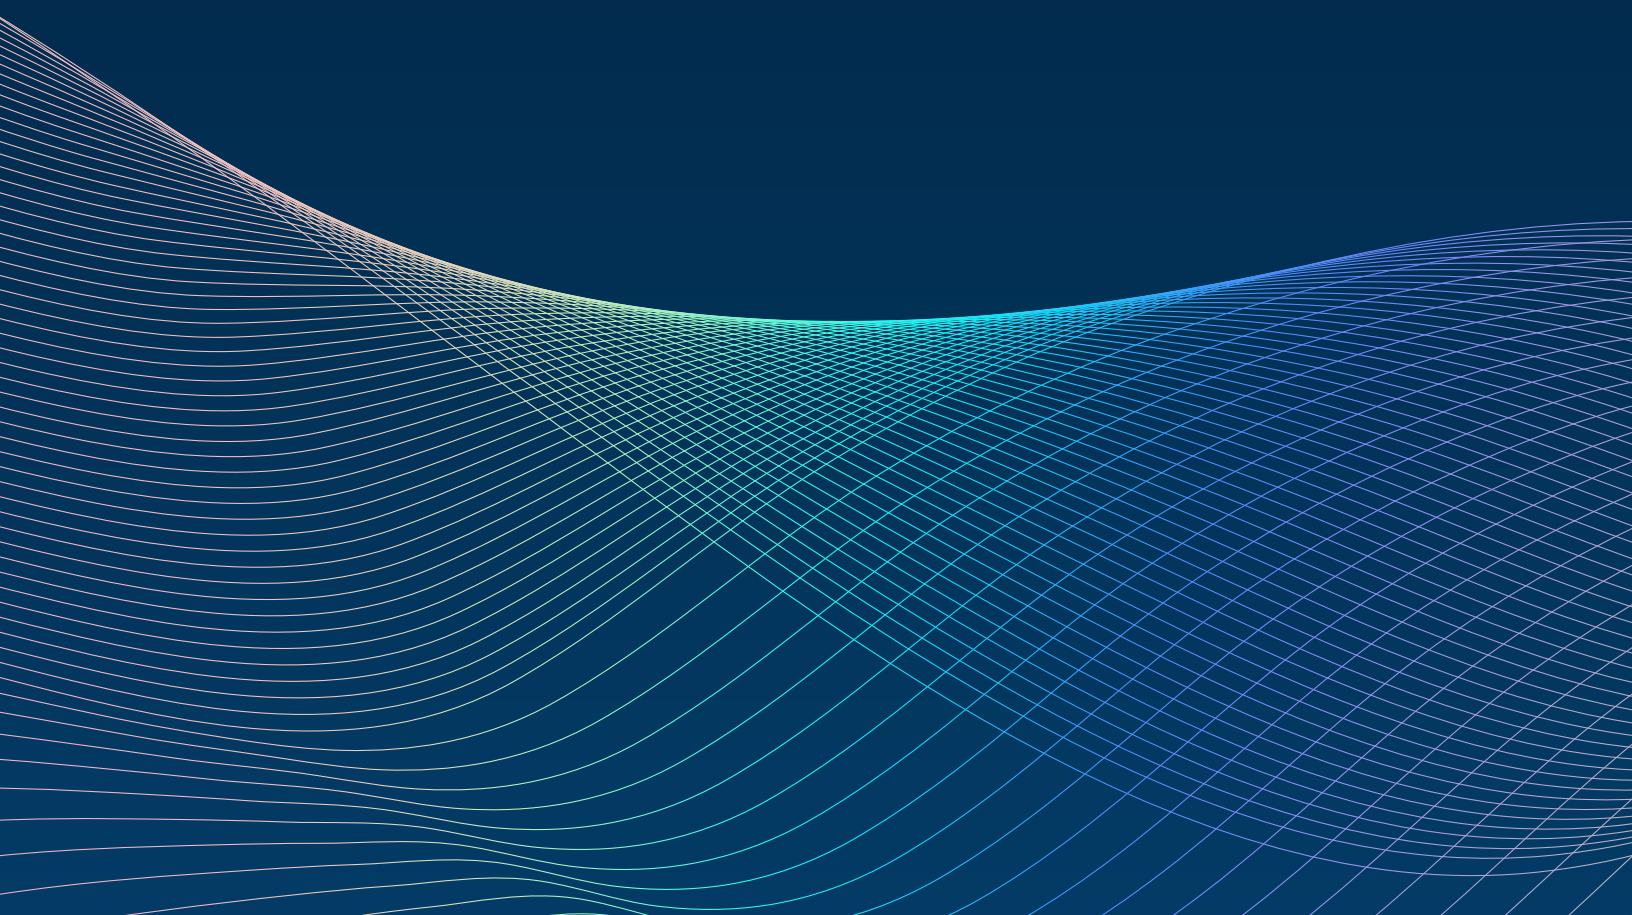

Introduction

Every major shift in the internet has forced a rethink of how trust is established online. From the emergence of ecommerce on the early web to the shift to mobile and social media platforms, each digital wave has both expanded business opportunities and created new openings for bad actors. In response, the industry has repeatedly introduced novel identity and fraud-prevention techniques in an attempt to keep pace.

Another such moment is upon us now. As artificial intelligence reshapes how people interact online, the reach and sophistication of fraud are seeing dramatic acceleration. Generative AI has lowered the barrier to creating fake personas, falsifying documents, and impersonating real people at scale. As a result, fraud losses are projected to reach **$40 billion USD** globally within the next few years, driven in large part by AI-enabled attacks.

In the new AI-first internet, the data sources historically used to verify identity and fight fraud no longer work. Signals that were once effective—static identity attributes, knowledge-based questions, even some biometric checks—can now be spoofed, purchased, or synthesized. To keep up, the industry needs a new foundation for trust.

Network insights, or more specifically, each person's unique behavioral footprint across the broader financial and app ecosystem, offer a solution. Grounded in long-standing relationships and individual patterns, this activity is difficult to fake and deeply connected to fraud outcomes. When viewed across institutions and platforms, it provides a resilient, high-signal lens into intent, risk, and legitimacy.

To better understand why financial network insights have emerged as the next foundation for trust, it's worth examining how each past wave of digital innovation has demanded a new approach to trust and fraud prevention.

01 The evolution of identity data and fraud protection

Each major innovation in identity verification over the past thirty years has emerged in response to a broader societal or technological shift.

The dot-com boom: As online commerce took off, identity verification leaned on **credit bureau data** (including names, addresses, and dates of birth) to establish who someone was.

The growth of the internet: As fraud followed users online, **knowledge-based authentication** emerged, asking questions derived from personal history (such as the name of a first pet or a favorite singer).

The rise of the gig economy: Marketplaces needed to verify real people quickly, leading to widespread adoption of **government ID document checks.**

The crypto take off: Fully digital financial services created demand for **liveness verification**, ensuring a real person was present at onboarding.

At-scale social media: Behavioral analytics, such as typing cadence and device behavior, added a layer of continuous fraud risk assessment.

Each of these approaches helped solve the problem of the given moment. Fraudsters, however, adapted in kind, weakening defenses. Today, many of these signals are widely available, increasingly commoditized, or vulnerable to AI-driven manipulation when used in isolation.

Notably, it has been nearly a decade since a truly new data source for identity and fraud prevention entered the market, despite fraud evolving faster than ever.

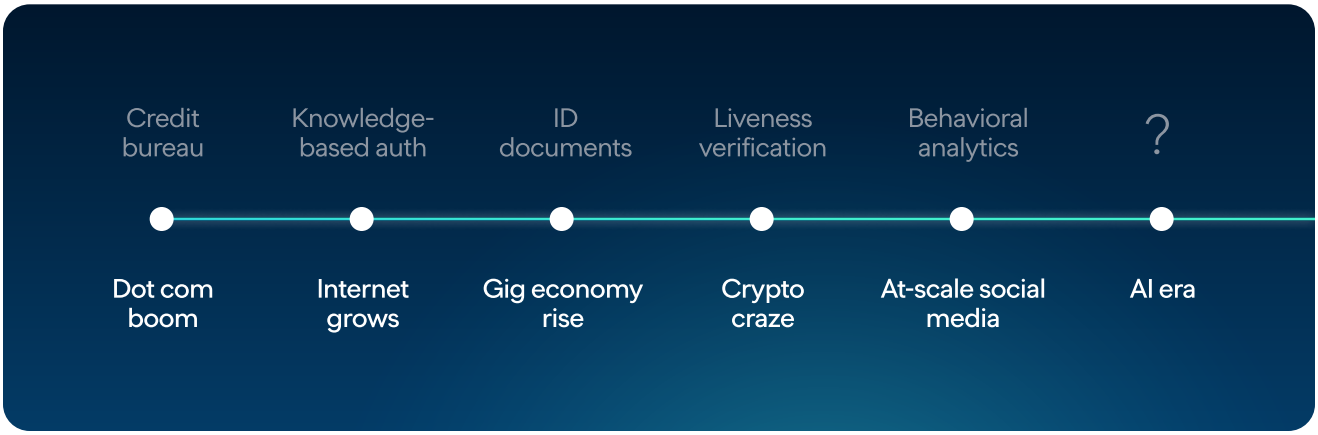

02 The AI era: reshaping how we transact online

> AI driven fraud is exposing the limits of identity controls that were designed for point in time verification rather than continuous assurance. As digital interactions become more autonomous and interconnected, institutions need signals that persist across the lifecycle and hold up at scale.

Sam Abadir

Research Director for Risk, Financial Crime, & Compliance,IDC

Artificial intelligence fundamentally changes digital interaction, particularly in financial services. Automation, personalization, and intelligent agents create more seamless experiences, but also new avenues for abuse. As AI reshapes business and online interactions, new questions are emerging for businesses about who they're transacting with and how they're transacting.

For example, in *agentic commerce,* AI agents act on behalf of users, organizations must verify not only the individual, but the relationship between the person and the agent operating in their name. Is the agent authorized to make X purchase? Is it authorized to purchase Y item? Is the purchased quantity correct? Businesses will need to make sense of these questions and many others,

In this new reality, contextual signals unique to each user emerge as a key and continuous means of authentication. Ongoing identity verification has thus become table stakes: accounts must be protected throughout their lifecycle, as fraud grows harder to detect due to the weaponization of AI. Organizations must verify every single interaction a user has with their business—that means authenticating every login, every click of a button, and every transaction to proactively identify account takeovers and first party fraud.

The result is a compounding challenge in which organizations must move faster, reduce friction, and personalize experiences, all while distinguishing legitimate users from increasingly sophisticated fraudsters.

Identity, fraud, and AI: why old methods fall short

AI has amplified several longstanding weaknesses in traditional identity and fraud approaches, including:

> Advances in AI have made it increasingly difficult to trust surface-level signals, making human-only detection ineffective for identity verification and fraud prevention. I've seen firsthand how static data and risk signals viewed in isolation lead to missed fraud.
>
> Taking a proactive approach that connects signals, applies context, and adapts to evolving fraud tactics allows teams to identify potential red flags earlier, before fraud occurs.

Jennifer Pitt
Senior Fraud Analyst,
Javelin Strategy & Research

Widespread availability of personal data: Breaches and data markets have made static identity information easy to obtain.

Synthetic identity creation: AI enables the fabrication of convincing personas, even where no real data previously existed—powering scams such as fake rental listings or loan fraud.

Deepfake technology: Audio, video, and image manipulation are challenging existing document checks and biometric verification.

Location and device spoofing: Fraudsters can mask their true environment, undermining geolocation and device-based controls.

The result is an environment where organizations struggle to trust who is showing up at the digital front door, whether a transaction or action is legitimate, and whether the signals they rely on reflect reality. To compensate, many teams overcorrect by adding friction, introducing manual reviews, or by requiring more user input. This slows growth, frustrates customers, and still fails to stop determined fraudsters.

The industry is overdue for its next wave of fraud-fighting innovation. The question is not whether change is needed, but what unique combination of data, insights, and analytics can meet this moment.

03 The importance of a digital financial footprint

Whatever the method or attack vector, fraud is always a matter of finances. That means the best way to identify and fight it is by following a person's financial footprint.

> When viewed at a network level, financial data shifts identity from a static check to an operational control that can be monitored, evidenced, and trusted over time.

Sam Abadir

Research Director for Risk, Financial Crime, & Compliance,IDC

A financial footprint reflects the real ways in which an individual behaves and transacts across the broader digital financial ecosystem over time. Combining device and IP information, or remembered user's interactions with apps and services, with transactional patterns like large amounts in and out of an account or transacted between their accounts can provide a new signal for fraud. Unlike static attributes or surface-level behaviors, it captures real-world patterns that are hard to fabricate.

In practice, a typical sequence might have a certain number of credit transactions, a certain number of merchants paid in a given time period, and a certain number of transfers out of a savings account. But, if behavior suddenly changes and the user starts signing up for multiple apps they've never used before at a fast velocity, or sending large sums via peer-to-peer apps, it could be a sign of fraud.

Crucially, financial data becomes exponentially more powerful when viewed across networks, not just within a single organization.

Consider a simplified example, in which a user signs up for a new app and passes identity checks without issue. Everything looks normal in isolation, but by tapping into a broader network view, we see that:

- The same phone number was recently used to open a crypto wallet elsewhere.
- This newly linked account was connected by a device previously flagged for promotional abuse on a peer-to-peer payments platform.
- The linked bank account has cycled rapidly between multiple fintech apps in the past week.

By analyzing different device attributes and its velocity across accounts, organizations can get a more holistic view of fraud risk. Individually, none of these signals are definitive. Together, however, they form a clear picture of multi-platform abuse—behavior that is nearly impossible to detect from within a single set of four walls.

This is the core advantage of financial network intelligence: it reveals patterns of intent across the ecosystem.

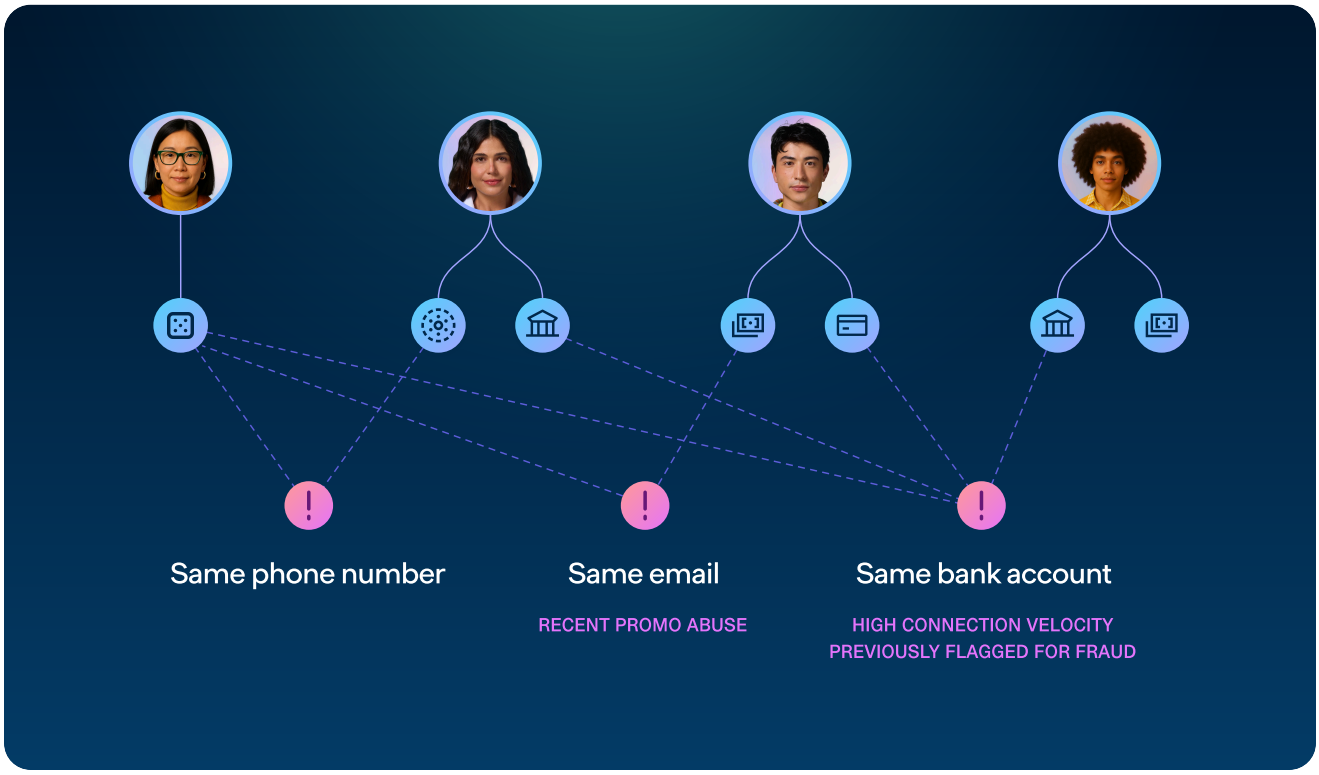

04 Financial network intelligence: better suited for the AI era

> For too long, fraud fighters have worked in silos, limiting their ability to see the full picture of each customer. A holistic view of each customer, **both within and across organizations,** enables **better risk assessment** and **helps uncover patterns and fraud rings** that would otherwise remain hidden.

Jennifer Pitt

Senior Fraud Analyst,
Javelin Strategy & Research

Understanding an individual's financial footprint offers three properties that ensure their unique resilience against fraud in an AI-first world:

01 **Verified data**

Modern financial data is consumer-permissioned and grounded in real banking relationships. It reflects actual usage with real bank connections, not easily purchased records.

02 **Cross-institutional activity**

Fraud never occurs in isolation. Financial activity spans banks, fintech apps, marketplaces, crypto exchanges, and payment platforms. Seeing user behavior and devices across a diverse range of organizations in a privacy-friendly way is essential to understanding true risk.

03 **Network intelligence**

Patterns emerge across a broad and diverse network. On the one hand, consistency helps organizations recognize and re-recognize legitimate users. On the other hand, anomalies like unusual transaction velocity, abrupt behavioral shifts, or unexpected account connections signal potentially illegitimate activity.

Recognizing trusted users is just as important as identifying fraudsters. Cross-network intelligence enables both.

05 Plaid's role in the next wave of fraud prevention

No single organization can see the full web of financial relationships on its own. But platforms that sit at the center of the financial ecosystem can come close.

Through its role powering digital financial services, Plaid has connected hundreds of millions of people with financial accounts across tens of thousands of institutions to thousands of apps and services. This position enables a unique form of network intelligence that reflects how users, devices, and bank accounts interact across the ecosystem.

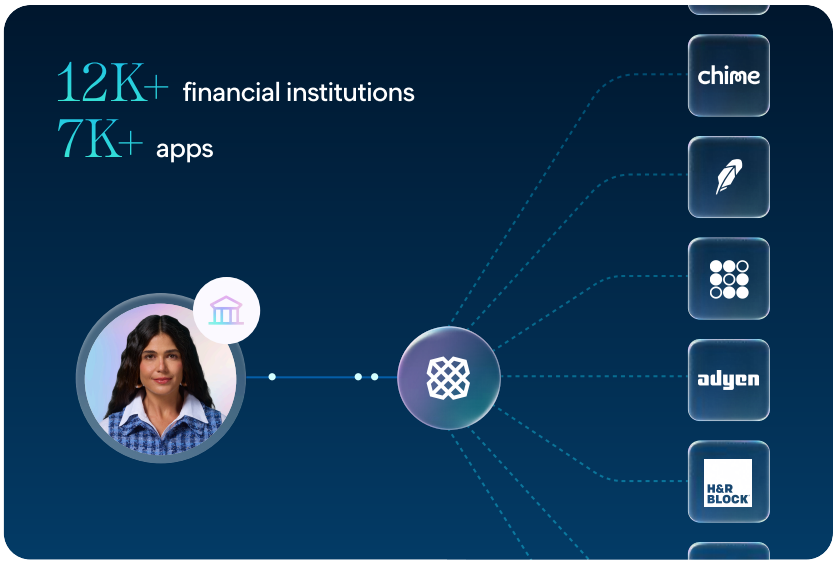

Whether the goal is stealing or laundering money, fraudulent activity leaves traces in the financial system. Following those traces—across platforms and over time—provides the most durable foundation for fraud prevention in the AI era.

Plaid Protect applies this network intelligence to identity verification and fraud decisioning, helping organizations detect risk earlier, reduce false positives, and defend accounts throughout the user lifecycle. Importantly, it does so without introducing unnecessary friction for legitimate users.

40%

With targeted step-ups applied to just 5% of users, this approach could have surfaced 40% of first-party fraud that had previously gone undetected.

The value of network-level behavioral insights becomes clearest when applied to real-world fraud scenarios. Across a range of anonymized customer environments, financial footprints have consistently surfaced risk signals that were invisible within a single platform's view.

In one case, a publicly traded lending company was struggling to contain first-party fraud, a category that often evades conventional fraud controls. Applicants were passing identity checks and onboarding successfully, only to later default or abuse credit in ways that appeared legitimate on the surface. Existing consortium data and internal rules flagged only a small portion of this activity.

By incorporating cross-institutional behavioral signals—such as account behavior across apps, device relationships, and patterns of prior activity—risk teams were able to identify potential fraudsters earlier in the funnel. **With targeted step-ups applied to just 5% of users, this approach could have surfaced 40% of first-party fraud that had previously gone undetected.** Importantly, the majority of legitimate users would experience no added friction.

47%

Stepping up identity checks on only one in ten users would have caught 47% of the fraud in question.

In another case study, a digital financial services company applied network-based financial intelligence to its most challenging fraud cases: accounts that had already passed sophisticated KYC and transaction monitoring controls, but were later closed for fraud. Based on Plaid-identified device signals alone, the study showed that **stepping up identity checks on only one in ten users would have caught 47% of the fraud in question.**

These results and others point to a broader conclusion: when behavioral insights are leveraged on a network level rather than a point-in-time signal, they enable a fundamentally different approach to fraud prevention—one that is both more effective and more user-friendly in an AI-driven threat landscape.

06 Build stronger defenses now

Now is a mission-critical time to look beyond traditional methods of fraud mitigation to carve out durable fraud prevention for your business. Meeting this moment requires more than incremental improvements to existing controls; it requires combining the same technologies fraudsters are using with data sources that are significantly harder to fake or manipulate.

Use AI to fight AI (for good): Applied thoughtfully, AI can strengthen fraud defenses rather than undermine them. Machine learning models can identify subtle, emerging fraud patterns that static rules and manual review often miss, particularly when trained on high-quality, real-world signals. From an operational perspective, natural language tools can accelerate investigation, surfacing trends across large datasets and enabling faster iteration on rules and risk strategies.

Pair AI with cross-institutional network intelligence: The effectiveness of AI tools depends on the data they're trained on. AI amplifies signal quality just as readily as it amplifies noise. By incorporating trusted, consumer-permissioned banking data and network-level insights into fraud models and rules, organizations can move beyond point-in-time verification to a more complete view of fraud risk. Patterns that appear benign within a single platform, become clear when viewed across the broader ecosystem.

Together, AI and network intelligence enable more precise decisioning: identifying high-risk activity earlier, stepping up only where needed, and preserving low-friction experiences for legitimate users. Moreover, by grounding fraud prevention in the reality of user behavior across the digital financial ecosystem, you can grow your defenses with the internet's next wave.

Plaid Protect

Plaid Protect brings fraud-fighting capabilities together in a single fraud decisioning platform. Built on Plaid's network of over one billion device connections to hundreds of millions of connected accounts, Protect combines machine learning models with cross-institutional insights to help you detect fraud throughout a user's lifecycle. By grounding AI-driven decisioning in real financial behavior, Protect enables you to catch fraud earlier and more often, while applying friction selectively to maintain seamless customer experiences. As financial services continue to evolve, so does Plaid's network intelligence—helping you build safer, more resilient products at scale.

Learn more