WHITEPAPER

Building Trust in Agentic Commerce

sardine PayOS

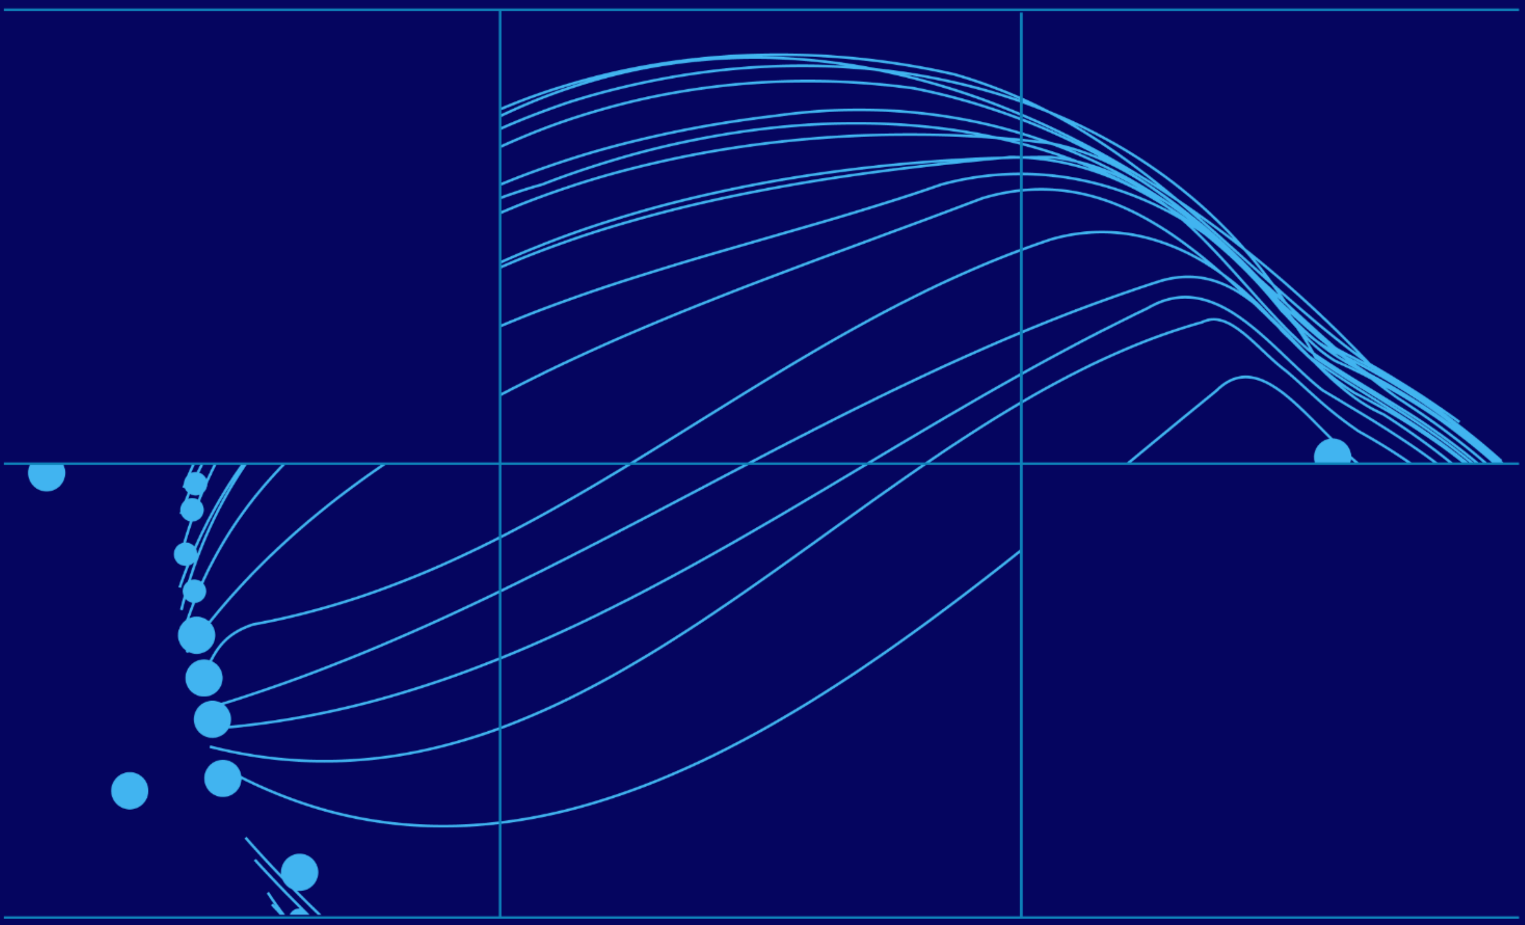

SARDINE.AI PAYOS.AI 2025

Table of Contents

Overview

When AI Agents transact, commerce shifts from a user making clicks on a screen to agents following a human's (or business') **intent**. AI Agents are now embedded in browsers, apps, and LLM chatbots and can make payments. This volume is set to skyrocket, raising fundamental new questions about trust, authorization, and fraud.

The payments industry has started to react with multiple initiatives emerging, designed to tackle these challenges.

- Visa has introduced Visa Intelligent Commerce[1], and Mastercard has launched Mastercard Agent Pay[2], each signaling a publicly stated commitment to building native infrastructure for agent-led transactions. Visa, in particular, has publicly stated its emphasis on credential-level security, including token provisioning, passkey authentication, and secure agent credential management.
- Companies like PayPal and Coinbase have publicly stated their belief in agentic commerce by rolling out developer-focused tools such as the PayPal Agent Toolkit[3] and Coinbase AgentKit[4], underscoring a shared vision of a future where intelligent agents can transact securely and seamlessly on behalf of users.

This shift introduces an opportunity for merchants to address fraud and establish trust earlier in the shopper journey at the point of intent. Since agentic commerce moves the communication with humans upstream from traditional checkout flows, it is essential to strengthen fraud prevention and trust layers accordingly. The challenge is ensuring legitimate AI agents aren't incorrectly flagged as malicious bots. This paper details our framework for trusted AI agent transactions, from intent to payment completion, with Sardine and PayOS.

The framework moves verifying agents, managing consent, and fraud signals to the point of intent, ensuring trust is established before any downstream transaction steps. Moving the risk and trust assurance to the point of intent ensures trust and safety are built into the earliest stages of agentic commerce, **aligning security with the moment a transaction is conceived, not just when it is executed.**

PayOS and Sardine have partnered to create a framework for trusted agentic commerce pathways. This framework enables legitimate AI agents to complete merchant checkouts while blocking malicious actors through global, real-time, design-level security measures.

[1] Visa Intelligent Commerce, https://corporate.visa.com/en/products/intelligent-commerce.html

[2] MasterCard Agent Pay, https://investor.mastercard.com/investor-news/investor-news-details/2025/Mastercard-Unveils-Agent-Pay-Pioneering-Agentic-Payments-Technology-to-Power-Commerce-in-the-Age-of-AI/default.aspx

[3] PayPal Agent Toolkit, https://developer.paypal.com/community/blog/paypal-agentic-ai-toolkit/

[4] Coinbase Agentkit, https://www.coinbase.com/developer-platform/discover/launches/introducing-agentkit

The paper outlines specific roles, responsibilities, and risk management best practices, with an open invitation for industry collaboration and feedback. In the whitepaper, we introduce a Trusted Agent Directory that helps merchants mitigate the information lost during an Agent-led checkout. The Trusted Agent Directory provides merchants with information that consumers provide at the Point of Intent, and allows merchants to distinguish between authorized Agents and malicious bots.

Definitions and Roles in Agentic Commerce

Agent

AI-driven software capable of performing tasks such as:

- Browsing product listings
- Filling out shipping and billing forms
- Collecting or retrieving stored user payment credentials
- Executing payments at the point of sale

Agents may be embedded in:

- Browsers or extensions (e.g., shopping companions, coupon finders)
- Mobile or desktop apps (e.g., personal finance or inventory managers)
- Creator tools (e.g., livestream overlays or affiliate helpers)
- Smart devices (e.g., AR glasses, voice assistants)

Agents can act synchronously (when triggered by user interaction) or asynchronously (autonomously polling needs and acting without immediate user involvement). They are capable of hopping between domains, managing form autofill, and programmatically submitting payment requests on behalf of the user.

Agents are not just browsing-helpers, they are transaction-executors.

Brand

The destination where a purchase is completed. Brands authorize or deny payment attempts, relying on payment credentials and context provided by upstream sources.

PayOS (Payment Infrastructure for Agentic Commerce)

A global agentic payments and billing platform that provides:

- Know Your Agent (KYA) and payment credential distribution: Secure compliance check of agentic developers or businesses, linked to verified developer identities, and distribution of payment credentials.
- Consent orchestration: A dynamic, revocable consent management layer where users can configure agent permissions by transaction amount, domain, frequency, or context.
- Fraud data sharing: With user consent, generates and standardizes structured metadata, point of intent information, that travels with each intent, including agent identity, consent method, timestamp, device ID, and origin URL.
- Trusted transaction routing: With proper user consent, ensure validated agent payment request is routed to brand checkouts, processors, and fraud engines downstream, reducing exposure to spoofed or rogue agents.
- Step-up decision support: Triggers additional verification (e.g., biometric or OTP) when risk scores cross brand-defined thresholds.
- User dashboard and personal finance control: PayOS offers a full-featured dashboard where users can:
    - View agent activity and transaction logs
    - Enable or disable individual agents at any time
    - Set consent levels (manual approval, auto-consent, amount thresholds)
    - Adjust “human-in-the-loop” preferences for different agents or domains
    - Manage recurring purchases and payment limits
    - Monitor agent-initiated spending through a built-in personal finance manager

This gives users both transparency and control, ensuring that agentic commerce happens with them, not around them.

Point of Intent

The Point of Intent is the true origin of a user’s purchase decision. It is where interest is sparked and buying momentum begins, often far upstream from the brand’s website or checkout flow (Point of Sale).

These can include:

- **Blogs, videos, livestreams:** A review for a sneakers on a video streaming platform, or a social media influencer showcasing a product can initiate purchase interest instantly.

- **Affiliate links and ads:** Embedded links in newsletters, price comparison engines, or social media stories drive users toward specific products or promotions.
- **Marketplaces and aggregators:** Platforms that collect intent across multiple brands before handing users off to a final transaction endpoint (point of sale).
- **Voice assistants and smart speakers:** Commands like “Buy more dog food” or “Order running shoes” bypass traditional shopping flows entirely.
- **Augmented reality and vision interfaces:** Pointing your phone at a lamp and asking “Where can I buy this?” creates intent without typing or browsing.
- **Agentic interfaces:** Browser extensions, chatbots, or personal finance agents that monitor user needs and proactively initiate purchases.
- **Social and conversational commerce:** In-app shopping flows in platforms where the checkout is triggered within a conversation.

What these all have in common is that they decouple discovery (point of intent) from the final purchase location (point of sale). The agent, acting on behalf of the user, takes over the journey from here, navigating across domains, retrieving credentials, and executing the purchase autonomously.

Modern commerce begins where intent is expressed, not where checkout happens.

Sardine (Fraud Intelligence Partner)

Performs:

- **Real-time risk scoring:** Calculates transaction risk using behavioral biometrics, device intelligence, and contextual information.
- Session and device analysis: Tracks how users (and agents) interact from the point of intent through to checkout, identifying signs of scripted automation or session hijacking.
- **Anomaly detection:** Uses machine learning to flag transactions that deviate from the expected pattern of a user, agent, or environment.
- **Fraud intelligence at the Point of Intent:** Gathers behavioral and environmental context (e.g., click patterns, OS fingerprints, language mismatches) at the earliest stage (well before checkout) to proactively identify malicious automation.

- **Fraud Tool at the Brand:** Brand-side decision engine that decides whether to authorize a transaction based on shared information and risk scoring.
- **Collaborative decisioning with PayOS:** When Sardine identifies elevated risk or insufficient data, it can work with PayOS to initiate additional step-up challenges or user outreach to improve decision confidence.

Trusted Agent

A Trusted Agent is one that:

- Is formally registered in our Trusted Agent Directory
- Is linked to a verified developer identity or organization
- Has a trackable operational history (incl. transaction metadata, error rates, risk events)
- Has a dynamic trust score, regularly updated based on behavior and fraud information
- Has completed security/compliance audits to ensure safe handling of user credentials
- Is only granted scoped access to user payment methods based on granular consent rules (amount limits, frequency, merchant domains)

Trusted Agents are issued signed credentials by PayOS that must be presented at the time of transaction. These credentials are invalidated if the agent violates trust policies or if the user revokes consent.

Trust isn't static. Agents must continuously earn and maintain it through secure, reliable, and transparent operation.

User

The consumer who owns the payment method and authorizes agents to act on their behalf, with full control over consent preferences.

The Global Challenge: Trust in an Agentic Future

Commerce is becoming decentralized, contextual, and autonomous. Historically, users discovered products through creator-led content, affiliate links, or search, but the final purchase still required them to click through and manually navigate to a brand's site to complete the transaction. Human-driven intent was converted into action through redirection, form-filling, and checkout flows owned by the brand.

That model is now being upended.

Transactions can now be initiated from any Point of Intent such as:

- Creator-led experiences (e.g., social media content, creator sharing platforms, blogs)
- Affiliate content and influencer links
- Voice and vision-enabled agents (e.g., in phones, AR glasses, assistants)
- Embedded AI agents in browsers or apps

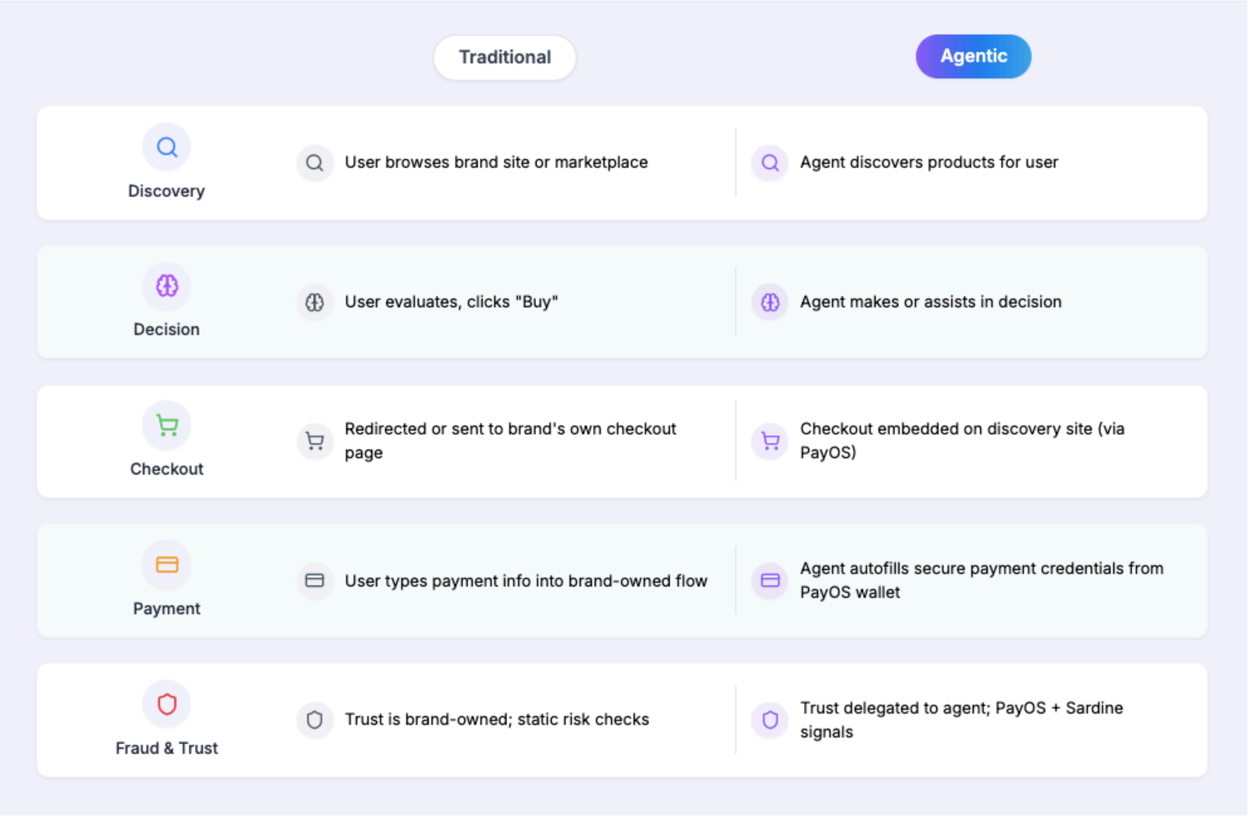

	Traditional	Agentic
Discovery	User browses brand site or marketplace	Agent discovers products for user
Decision	User evaluates, clicks "Buy"	Agent makes or assists in decision
Checkout	Redirected or sent to brand's own checkout page	Checkout embedded on discovery site (via PayOS)
Payment	User types payment info into brand-owned flow	Agent autofills secure payment credentials from PayOS wallet
Fraud & Trust	Trust is brand-owned; static risk checks	Trust delegated to agent; PayOS + Sardine signals

The Evolution of Checkout Experiences: Traditional vs Agentic Checkout Flow

These Points of Intent generate a purchase decision before the user ever reaches a brand's checkout page, if they reach it at all. Agents act on this intent by taking the user's payment credentials, navigating to brand sites, and executing purchases autonomously.

This is not theoretical. It is already happening, and it is global in nature.

But as automation rises, so do fraud vectors. Malicious bots, impersonators, and rogue agents can exploit open commerce endpoints unless a trust architecture is in place.

Brands face a stark new set of challenges:

- How do we know this agent is real?

- Is the truly delegated to make a purchase on behalf of the user?
- Has the user consented to the transaction?
- Can we trust the payment credential they are presenting?
- Or is our fraud system going to decline this potentially good agent representing a legitimate transaction because it thinks it's a malicious bot?

Without a universal system of agent identity, user consent, and real-time fraud intelligence, commerce is headed toward a fragmentation of trust, broken checkout experiences, and lost sales for merchants.

Why Traditional Fraud Tools Aren't Enough

Fraud prevention tools today are built for a world where humans are the actors. These systems rely heavily on behavioral cues that emerge during manual checkout experiences like mouse movement, keystroke dynamics, session duration, or mobile tilt.

Typical fraud engines are:

- **Calibrated to detect unusual human behavior:** They look for inconsistent behavior patterns that deviate from the historical norm of a user.

- **Triggered by real-time anomalies:** Such as erratic mouse movements, unfamiliar device fingerprints, mismatched geo-IP data, or unusual login sequences.

- **Context-bound to a single website:** These tools operate within the four walls of a merchant's environment, lacking upstream visibility into where or how a transaction was initiated.

But agentic commerce breaks all of these assumptions:

- **The user is absent at checkout:** There's no human cursor to track or button click to analyze. It may be a piece of software navigating through.
- **The agent is the actor:** The entity completing the transaction is autonomous software, often working in the background or across multiple platforms.
- **Session boundaries dissolve:** An agent might start on a blog, pull payment credentials from a secure vault, and land on a brand's checkout, all in milliseconds.
- **Risk must be evaluated before transaction execution:** By the time a transaction reaches the brand, it may already be too late to detect fraud if upstream context hasn't been captured.

In this new reality, traditional merchant-side fraud tools have limited visibility into upstream intent. They cannot see:

- Who the agent is
- Whether the agent has been approved by the user
- If consent was given for this transaction
- The full path from intent to execution

Take returns fraud as a case in point. Retailers are experiencing a surge in return abuse and refund fraud, where malicious actors exploit generous return policies, for example, by sending back only part of an order or substituting items with counterfeit goods.

In the world of Agentic Commerce, traditional chargeback and dispute resolution processes must be reimagined. Consider common return reason codes like Item Not Received (INR)[5] or "Significantly Not As Described (SNAD)[6]." How should these outcomes affect the trust score of both the AI agent and the end user it represents?

We believe trust should be transitive: if an agent consistently enables purchases that lead to legitimate disputes or excessive returns, that behavior should reflect on the user's trust profile as well. But for this model to work, merchants must have clear visibility into the user behind the agent.

This is where our proposed Trusted Agent Directory plays a pivotal role. Not only for purchase authorization, but also for managing post-transaction workflows like returns and chargebacks.

Without deep, interoperable connections between agent identity, user consent, and transaction context, current fraud systems will struggle. Either by rejecting legitimate transactions or by failing to detect fraud at scale.

The future of fraud prevention is at the point of intent.

That's why platforms like PayOS and Sardine are collaborating to embed trust earlier in the journey, before the transaction even begins.

[5] ¹ INR is a chargeback reason code used when a customer claims non-receipt of goods or services. Visa uses Reason Code 13.1; Mastercard uses 4855 or 4853.

[6] ² SNAD is a chargeback reason code used when a customer claims the item received is materially different from what was advertised. Visa uses Reason Code 13.3; Mastercard uses 4853 in relevant cases.

The PayOS & Sardine Solution: Infrastructure for Secure Agentic Transactions

1. Trusted Agent Directory

PayOS maintains a dynamic and trusted directory of registered agents to ensure only verified, compliant agents can receive payment credentials.

This living registry includes, but not limited to, the following attributes:

- **Agent Identifier:** A unique ID assigned to each agent.
- **Linked Developer ID:** The verified ID of the developer or organization responsible for the agent.
- **Platform:** The platform where the agent operates (e.g., Chrome extension, iOS app).
- **Trust Score:** A dynamic score reflecting the agent's performance and information of trust and safety.
- **Revoked Status:** An indicator showing whether the agent is currently suspended or revoked.
- **Permitted Domains:** A list of domains the agent is allowed to operate on or interact with.
- **Last Audit Date:** Timestamp indicating the last time the agent underwent a security or compliance review.

This registry is designed for interoperability across the ecosystem and is open to other agentic commerce providers and third-party fraud platforms. Participants can join the Trusted Agent Registry via a Know Your Business (KYB) process, where they are issued a private-public key pair. Other participants can use the public key to validate the participant's identity.

When an agentic commerce provider registers an agent on behalf of a customer, it signs the intent attributes with its own private key. Fraud providers can then verify this signature by checking it against the agentic commerce provider's public key in the registry.

This registry is shared with third-party trust and risk platforms, as well as large merchants, enabling them to incorporate verified agent data into their risk-based decision-making process at the time of payment processing.

2. Point of Intent (POI) Information

Every transaction initiated by an agent is accompanied by Point of Intent Information, a rich set of standardized metadata that travels from the point of intent through to the brand's checkout and fraud systems. This information package enables precise verification, real-time risk assessment, and fine-grained authorization decisions.

Key metadata elements may include, but not limited to:

- **User Consent Method and Timestamp:** Details on how and when the user granted permission for the agent to act, whether through explicit approval (e.g., manual confirmation), pre-approved scopes (e.g., limited amount or merchant domains), or always-on consent for trusted agents. Timestamping provides an audit trail to verify consent freshness.
- **Agent Identity and Origin:** Unique identifiers for the agent, including its registered ID in the PayOS Trusted Agent Directory, developer credentials, platform (browser extension, mobile app), and the originating environment (e.g., URL or app). This helps brands confirm the legitimacy of the actor completing the transaction.
- **Brand Domain and Transaction Context:** The exact merchant domain or brand URL where the purchase is executed, including transaction amount, currency, item or service description, and timestamp. This allows contextual risk scoring based on merchant history and transaction patterns.
- **Delegation level:** Defines the scope of agent authority granted by the user, such as:
    - Manual: User approval required per transaction
    - Pre-approved: Scoped permissions for specific merchants, amounts, or frequencies
    - Always-on: Autonomous agent actions within defined boundaries
- **Device and Session Information:** Metadata about the device from which the transaction is initiated, including device ID, IP address, geo-location data, browser fingerprint, and session duration. This aids in detecting anomalous device usage or session hijacking.
- **Consent Revocation Status:** Flags indicating whether the user has revoked or modified consent since issuance, enabling brands to reject or pause transactions from unauthorized agents.

- **Transaction History and Agent Trust Score:** Aggregated data on the agent's prior transaction success rate, fraud flags, and dynamic trust scoring, helping brands weigh transaction risk in real time.
- **Delegation Chain and Agent Hierarchy:** If an agent delegates action to sub-agents or plugins, the full chain of authority and related identities is included to maintain transparency and accountability.

This information is cryptographically signed by PayOS' Private key, standardized and shared downstream by PayOS to Sardine.

3. Fraud Intelligence Integration

PayOS works in close partnership with Sardine to deliver a multi-layered, proactive defense system that protects agentic commerce at every stage, from initial user intent through transaction authorization.

Key capabilities may include but not limited to:

- **Comprehensive Session and Device Analysis:** From the moment a user expresses intent to purchase, whether through content, affiliate links, or agent triggers, PayOS and Sardine analyze session behaviors and device fingerprints. This includes tracking mouse movements, click patterns, device hardware identifiers, IP reputation, and environmental information such as language settings and browser anomalies, enabling early detection of scripted automation or hijacked sessions.
- **Anomaly Detection Across Agents and Contexts:** Leveraging machine learning and historical transaction data, the system identifies unusual patterns not just in isolated transactions but across multiple agents, users, and contexts. For example, it flags agents suddenly accessing new merchant domains, performing atypical purchase volumes, or using unfamiliar devices.
- **Dynamic Risk Score Adjustment:** Real-time risk scores are continuously updated based on evolving behavioral information, agent trust scores from the Trusted Agent Directory, transaction metadata, and external threat intelligence feeds. This allows risk assessments to be contextually aware and adaptive.
- **Trusted Agent Directory Analysis:** Sardine integrates with PayOS' Trusted Agent Directory to cross-reference agent identity, ownership, historical performance, and revocation status. This ensures only agents with verified, good standing can proceed.

- **Rich Transaction Information Evaluation:** Prior to finalizing transaction processing, Sardine analyzes detailed transaction information from PayOS, including user consent parameters, delegation scopes, device information, and transaction context, to identify potential inconsistencies or elevated risks. Sardine also utilizes its own device intelligence information received from Point(s) of Intent.

- **Collaborative Step-Up and User Outreach:** When the combined intelligence detects ambiguous or high-risk scenarios, Sardine coordinates with PayOS to initiate additional friction only when necessary. This includes triggering step-up authentication (e.g., biometric verification, OTP challenges), deferring authorization for manual review, or conducting user outreach to verify intent and identity before processing the transaction.

- **Predictive Fraud Prevention:** By integrating diverse data sources and real-time information, the collaboration transforms fraud prevention from a reactive, post-transaction process into a predictive, front-line defense. This minimizes false positives, reduces unnecessary friction for legitimate users, and stops fraudulent activity before it impacts brands or consumers.

This tight integration ensures that the agentic commerce ecosystem remains secure, transparent, and trustworthy, enabling brands to confidently accept agent-initiated payments while safeguarding user autonomy.

Shift in Core Paradigms

The rise of agentic commerce represents a fundamental shift in how trust, control, and decision-making are distributed. As agents act on behalf of users, key paradigms like consent, fraud detection, and checkout ownership move upstream, enabling purchases to happen securely and instantly at the point of intent.

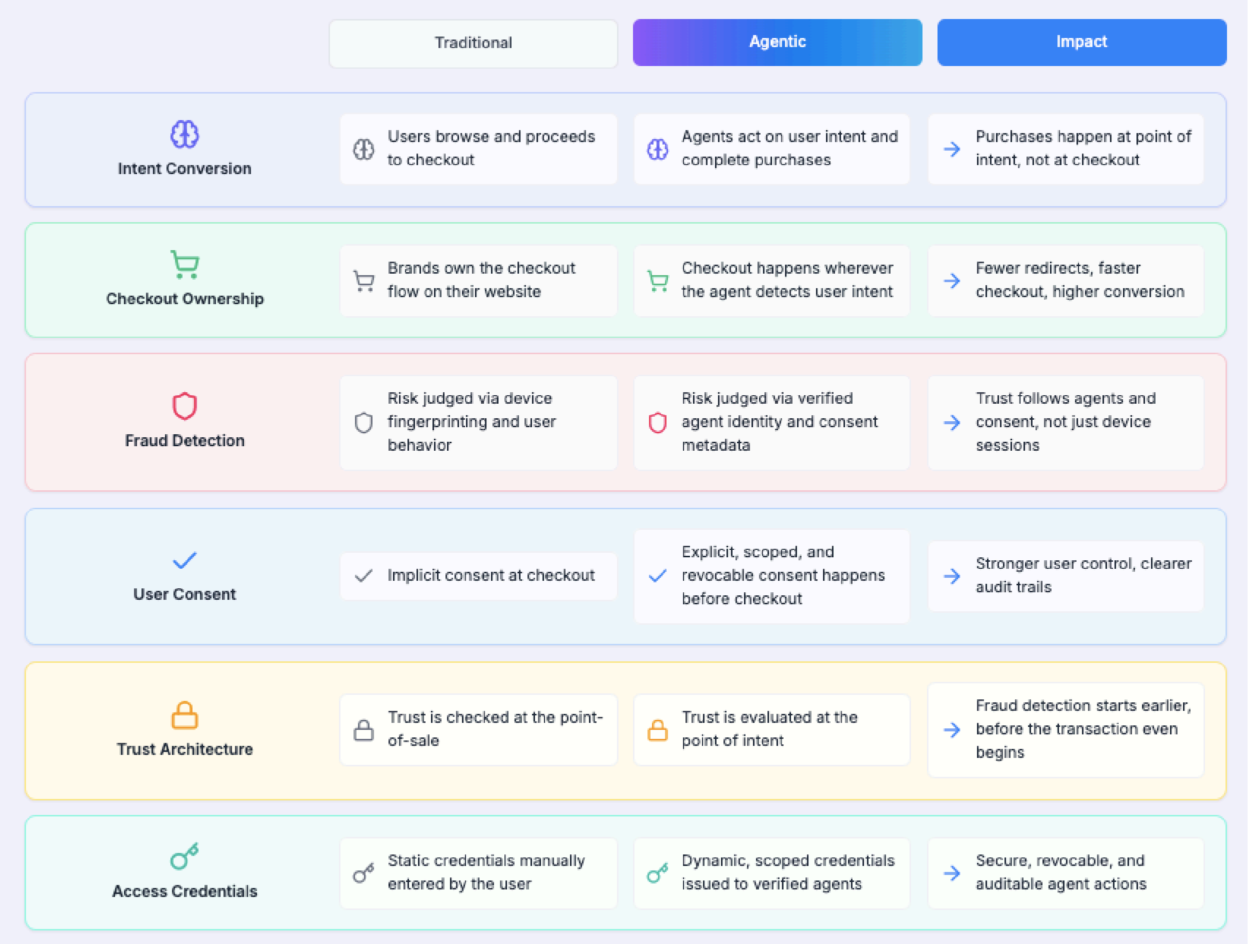

	Traditional	Agentic	Impact
Intent Conversion	Users browse and proceeds to checkout	Agents act on user intent and complete purchases	→ Purchases happen at point of intent, not at checkout
Checkout Ownership	Brands own the checkout flow on their website	Checkout happens wherever the agent detects user intent	→ Fewer redirects, faster checkout, higher conversion
Fraud Detection	Risk judged via device fingerprinting and user behavior	Risk judged via verified agent identity and consent metadata	→ Trust follows agents and consent, not just device sessions
User Consent	✓ Implicit consent at checkout	✓ Explicit, scoped, and revocable consent happens before checkout	→ Stronger user control, clearer audit trails
Trust Architecture	Trust is checked at the point-of-sale	Trust is evaluated at the point of intent	→ Fraud detection starts earlier, before the transaction even begins
Access Credentials	Static credentials manually entered by the user	Dynamic, scoped credentials issued to verified agents	→ Secure, revocable, and auditable agent actions

Shift in Core Paradigms: From Old to New

User-Centric Consent: Friction Where Needed, Autonomy Where Allowed

User trust cannot be sacrificed in the name of frictionless automation. Nor should every agent transaction be forced to a manual approval.

PayOS allows users various levels of permission-setting for agentic payments.

- **Human-in-the-loop (HITL):** User manually approves each transaction.
- **Pre-approval:** User grants an agent a limited scope (e.g., up to $200/month at X.com).
- **Always-on:** Trusted agents can act without intervention, within defined boundaries.

- **Custom/Hybrid:** Users can delegate pre-approval control for certain parameters (e.g., specific hours during the day when the human is away or during ‘do not disturb’ times, or up to certain transaction amounts) and require approval for all other cases.

Users can revoke access, adjust permissions, and view agent activity—all from a central dashboard.

Agents don’t just act. They act with permission.

Real-time Brand Decisioning (Joint Orchestration by PayOS & Sardine)

Sardine performs real-time risk scoring using behavioral intelligence, Point of Intelligence Information, and environmental data gathered from the full session, starting at the point of intent. These scores, combined with PayOS’ verified agent credentials and consent metadata, power a coordinated trust envelope delivered to the brand.

If uncertainty remains, the brand (or its processor) can dynamically request one of the following through PayOS, in collaboration with Sardine:

- Step-up authentication (e.g., biometric confirmation, OTP, email confirmation)
- Deferred authorization (place transaction on hold pending additional context)
- Agent denial (if risk score exceeds defined thresholds or agent is revoked)

This real-time loop between Sardine, PayOS, and the brand ensures that:

- Risk is addressed before transaction execution
- Friction is applied only when trust is in question
- User autonomy is respected, and consent remains central
- False positives are minimized, preserving conversion and user experience

Fraud prevention becomes a collaborative decisioning mechanism, not a siloed guess.

Global Implications

Fraud is borderless. In a world where agents operate across platforms, countries, and ecosystems, inconsistent trust standards become exploitable weak points. Without a shared infrastructure, bad actors will simply target the most vulnerable nodes in the network.

PayOS is building a globally interoperable, open ecosystem to address this fragmentation, in collaboration with Sardine, by:

- Standardizing agent registration through a universal Trusted Agent Directory that can be queried across participating platforms.
- Creating data-sharing protocols that ensure real-time transmission of agent identity, consent metadata, and behavioral context across the transaction path.
- Enabling fraud intelligence APIs where Sardine can access enriched session data, delegation history, and device information from PayOS to fine-tune risk assessments.
- Building feedback loops where Sardine flags anomalies or risk scores to PayOS, prompting dynamic updates to agent trust scores or revocation status.
- Collaboratively supporting required compliance by ensuring that user consent, credential usage, and step-up events are fully logged and auditable.
- Enriching global threat models by pooling anonymized fraud patterns across agents and regions, making it harder for localized fraud tactics to scale globally.
- Powering real-time orchestration at the point of checkout ensures that a transaction approved in one region by one fraud engine respects the global trust context shared between PayOS and Sardine.

Just as HTTPS became the global trust layer for web browsing, PayOS and Sardine are building the trust and security fabric for agentic commerce. One that is scalable, privacy-preserving, and collaborative by design.

User Story: Agent-Powered Purchase with Trusted Authorization

Problem:

Agent-initiated transactions are susceptible to getting blocked by merchant fraud detection systems because they resemble automated bot behavior: no mouse movement, unrecognized devices, and auto-populated form fields.

Solution:

PayOS and Sardine enable merchants to verify that agents are legitimate, users have provided proper consent, and transactions fall within approved limits through a Trusted Agent Registry combined with signed consent metadata. This system ensures trusted agents can complete purchases without being mistaken for fraud.

1. A user browsing Instagram sees a creator's shirt and tells their agent, "Find this same shirt in size M under $75."
2. The agent finds a matching item on a known brand's site and requests user approval.
3. The user approves the purchase, authorizing the agent to spend up to $75.
4. The agent initiates the checkout process and includes a signed package containing device ID, IP address, timestamp of consent, transaction amount, and wallet identity.
5. This package is signed using PayOS' private key to confirm authenticity.
6. During checkout, the merchant (or fraud provider) queries the PayOS Trusted Agent Registry for agent verification and transaction details. The agent then submits a session key along with relevant metadata as part of the transaction request.
7. Sardine uses the session key to look up the agent in the registry and verify the signature using PayOS' public key.
8. Sardine confirms that the request matches what was authorized, including device context and transaction scope.
9. If the verification is successful and risk is within acceptable limits, Sardine recommends transaction approval and the transaction proceeds for processing.
10. If the details don't match or the risk is high, Sardine suggests the merchant blocks the transaction or escalates based on fraud rules.
11. Only PayOS can sign registry entries, preventing unauthorized agents from spoofing valid ones.

Partnering to Build the Trusted Future of Agentic Commerce

If you are a:

- Brand aggregator building in agentic commerce
- Social network or affiliate marketer bridging product discovery and payment
- Merchant preparing for autonomous commerce
- Agent developer seeking an agentic payments solution and registry inclusion
- Acquirer building the next-gen payment layer

Then join us. Visit payos.ai and sardine.ai for more information.

Contributors

Johnathan McGowan is the cofounder and CEO of PayOS, a global platform for secure, card-native payment acceptance and billing in AI-driven commerce. He has over seven years of experience driving product and engineering initiatives in embedded payments, payment orchestration, and real estate fintech. Before PayOS, he contributed to the development of payment portals, SDKs, and terminals at TabaPay, he cofounded Trusty, a real estate payments and verification platform, and worked on VisaNet AI, fraud solutions, installment payments, and digital payment innovations at Visa. Known for bridging technical execution with strategic product vision, Johnathan consistently delivers solutions that advance payment security and developer enablement in complex ecosystems.

Aparna Krishnan Girish is the cofounder and Chief Product Officer of PayOS, a global platform for secure, card-native payment acceptance and billing in AI-driven commerce. She brings over two decades of experience across global money movement, mobile payments, QR, e-commerce tokenization, and risk tools. Previously, she was Head of Product at TabaPay and spent over a decade at Visa, where she played a key role in Secure Remote Commerce, payment tokenization, and has contributed to industry standards efforts across EMVCo, FIDO, and USPF, ISO, and W3C. Aparna holds a dozen patents in digital payments and has contributed to some of the industry's most widely adopted technologies.

Soups Ranjan, Ph.D., is the cofounder and CEO of Sardine, an AI-powered platform for fraud prevention and compliance. With over two decades of experience applying machine learning and artificial intelligence to combat fraud and financial crime, Soups has built a reputation as both a seasoned data scientist and a trusted executive. Prior to founding Sardine in 2020, he held senior leadership roles in fraud, compliance, and data at leading finance and technology companies including Yelp, Coinbase, and Revolut.

Zahid Shaikh is the cofounder and Chief Product Officer of Sardine. Before Sardine, Zahid was a product leader at Revolut, Uber, and PayPal. At PayPal, he created the Device Intelligence Product, which was responsible for $40 million per year in fraud reduction, and was recognized as the Top PayPal Inventor in 2016 with five approved patents for filing. Earlier in his career, Zahid held technical and engineering roles as a consultant at Verizon, Chase, and Hitachi. He graduated from the University of Mumbai with a Bachelor's Degree in Electronics Engineering and earned a Post-Graduate Diploma in Software Technology from the Centre for Development of Advanced Computing.

Simon Taylor is the Head of Strategy & Content at Sardine. With 20 years in financial services, today, Simon is trusted by CEOs, Regulators, and policymakers to explain the changes in finance and how they impact society. Simon started as a software engineer before working in cards, payments, correspondent banking, and ultimately becoming a consultant to the industry. Today, Simon runs fintechbrainfood.com, the weekly newsletter with over 40,000 subscribers, and is a regular in mainstream media outlets.

About Us

PayOS

PayOS is a next-generation global payments and billing platform powering agent-driven commerce. The platform enables agents to securely vault cards, streamline checkouts, send and receive payments, and manage billing, all through a unified, compliant system. With PayOS wallets, users link their card once for seamless payments across agent-driven workflows. Processor-agnostic by design, PayOS offers businesses secure, flexible integration with their preferred payment processors. By integrating agentic events, billing, and charging users, the platform enables agents to monetize users effectively. Learn more at PayOS.ai

Sardine

Sardine is the leading AI risk platform for fraud prevention, compliance, and credit underwriting, trusted by enterprises in over 70 countries. Using device intelligence, behavior biometrics, and machine learning, Sardine stops fraud in real time, streamlines compliance, and unifies data across risk teams. Backed by world-class investors including Andreessen Horowitz, Activant Capital, Visa, Experian, Moody’s, and FIS, Sardine is redefining risk management for the real-time economy. Learn more at www.sardine.ai