The Identity Layer for Agentic Commerce

Identity and delegation infrastructure for autonomous AI agents—IDX and KYA, the identity layer of Kachyng's Agentic Commerce Infrastructure

Kachyng Inc.

Part 1 of 8 · Patents & Patents Pending · April 2026

Executive Summary

AI agents are about to transact and act on behalf of humans and enterprises at a scale the current identity stack was not built for. McKinsey projects $3–5 trillion in autonomous commerce globally by 2030 and calls the missing authorization layer "a critical authorization failure below the surface of agent adoption." Today's identity infrastructure — OAuth, SAML, Active Directory, service accounts — was designed for a human logging into a service, not for autonomous software reasoning, delegating, and transacting across systems. This paper covers the identity layer of Kachyng's Agentic Commerce Infrastructure: IDX (Identity Exchange) and KYA (Know Your Agent). Orchestration (AGX) and settlement (PRX) are addressed in companion papers.

IDX and KYA together govern who an AI agent is, what it has been authorized to do, and whether it should be trusted at the moment of action. Kachyng has developed a patent-pending architecture for this layer that lets a human enroll once and delegate scoped authority to multiple agents while retaining oversight, revocation, and runtime risk control. The architecture addresses seven technology gaps through an integrated approach; orchestration (AGX) and processor execution (PRX) are treated in companion papers.

Seven Components of the IDX + KYA Identity Layer:

- **Agent Identity Framework:** Independent digital identity per AI agent, cryptographically bound to a human controller through Delegation Grant Tokens (DGTs), with ML-based trust scoring that adjusts each agent's autonomy over time
- **Trusted Gateway Architecture:** IP-based infrastructure control so compromised tokens cannot be used from unauthorized networks
- **Intelligent API Translation (Graphify):** Automatic conversion of legacy REST APIs to AI-optimized responses, with service-cooperative manifest discovery
- **Payment Scope in Delegation Grant Tokens:** Per-agent payment authority carried inside signed DGTs, spend caps, merchant allowlists, time windows, with biometric escalation on sensitive actions (execution mechanics covered in the PRX whitepaper)
- **Zero-Credential Login Agent:** Credentials stored in encrypted vaults accessible only from trusted infrastructure; never touch user devices, eliminating phishing and credential-theft attack surfaces
- **Rapid Revocation Network:** Publish-subscribe architecture for per-agent, per-class, or system-wide invalidation with cryptographic integrity and complete audit trails
- **Unified Adaptive Security:** Biometric-secured control surface with ML-based risk assessment, behavioral anomaly detection, and dynamic trust-score and permission adjustment (the KYA runtime trust layer)

Industry Overview

Market data as of Q1 2026. Sources cited inline; figures from third-party analyst reports.

Market Size

Kachyng is building Agentic Commerce Infrastructure for a category McKinsey projects will orchestrate $3–5 trillion globally by 2030 as AI agents shop, negotiate, and transact on behalf of humans and businesses. McKinsey describes the shift as "seismic," comparing it to the web and mobile revolutions and noting that payments, identity, and authorization are the enabling rails that will determine who captures demand as delegation expands.[1]

This whitepaper focuses on the identity-layer subset of that opportunity, IDX and KYA, which itself sits at the intersection of two large, fast-growing adjacencies. The global AI agents market was valued at approximately $7.6–8.0B in 2025 and is projected to reach $48–53B by 2030, representing a CAGR of 43–46%[2]. The adjacent digital identity solutions market stood at $44–47B in 2025 and is forecast to grow to $132B by 2031 at a 20% CAGR[3]. The narrower identity and access management (IAM) segment alone was estimated at $26B in 2025, growing to $43B by 2030 at a 10.4% CAGR[4]. Notably, the non-human identity (NHI) segment, machine identities for APIs, bots, and automated workflows, is emerging as one of the fastest-growing subsegments within digital identity, as enterprise machine identities now often outnumber human users. Payment processing, orchestration, and channel infrastructure sit outside this whitepaper's scope and are addressed in companion papers on PRX and AGX.

Growth Drivers

- **Enterprise AI agent adoption at scale:** Organizations across BFSI, healthcare, e-commerce, and IT are deploying autonomous agents for trading, customer service, procurement, and infrastructure management, creating urgent demand for identity and delegation infrastructure that existing IAM stacks were not designed to provide.
- **Proliferation of non-human identities:** Machine identities (API keys, service accounts, bot credentials) already outnumber human users in many enterprises. Managing, scoping, and revoking authority for these identities at scale is an unsolved infrastructure problem.

[1] McKinsey & Company, "The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants," October 17, 2025. Projects $3–5 trillion in orchestrated agentic commerce globally by 2030.
[2] MarketsandMarkets, AI Agents Market, 2025. BCC Research, AI Agents Market, January 2026. Grand View Research, AI Agents Market Report, 2025.
[3] MarketsandMarkets, Digital Identity Solutions Market, 2025. Grand View Research, Digital Identity Solutions Market Report, 2025.
[4] MarketsandMarkets, Identity and Access Management Market, 2025.

- **Credential-based attack surface:** Password and token compromise remains the dominant attack vector in enterprise breaches. Agent architectures that eliminate stored credentials and enforce infrastructure-origin trust represent a structural improvement over traditional authentication flows.
- **Multi-agent orchestration:** The shift from single-agent to multi-agent systems (projected 48.5% CAGR through 2030[5]) multiplies the need for hierarchical delegation, inter-agent trust scoring, and rapid cross-service revocation.
- **Regulatory pressure for AI accountability:** Emerging frameworks, NIST AI Agent Standards Initiative (Feb 2026), EU AI Act high-risk system requirements, and state-level AI governance laws, are creating compliance mandates around agent identity, auditability, and human oversight.

Regulatory Backdrop

The regulatory landscape for AI agent identity is crystallizing rapidly. In February 2026, NIST launched its AI Agent Standards Initiative through the Center for AI Standards and Innovation (CAISI)[6], focused on establishing interoperability and security standards for autonomous AI agents, including identity, authorization, and agent-to-agent communication protocols. NIST's January 2026 Request for Information on AI agent security[7] signals that formal federal guidelines are likely by 2027. In the EU, the AI Act's high-risk system provisions (effective August 2025, with compliance deadlines extended to 2027–2028 under the Digital Omnibus package[8]) impose transparency, human oversight, and audit requirements that apply directly to autonomous AI systems. At the U.S. state level, California and New York enacted frontier AI governance laws in late 2025[9], while a December 2025 executive order directed the creation of a national AI policy framework[10] to preempt a patchwork of state regulations. For companies building AI agent infrastructure, these converging mandates create both compliance requirements and a commercial opportunity: enterprises will need identity and delegation platforms that are audit-ready by design.

[5] MarketsandMarkets, AI Agents Market, 2025 (multi-agent systems segment).
[6] NIST, Announcing the AI Agent Standards Initiative, February 18, 2026.
[7] NIST, Request for Information Regarding Security Considerations for AI Agents, Federal Register, January 8, 2026 (docket NIST-2025-0035).
[8] EU AI Act (Regulation 2024/1689), effective August 2, 2025. EU Digital Omnibus package, late 2025.
[9] California SB-53 and New York RAISE Act, enacted late 2025.
[10] Executive Order, Ensuring a National Policy Framework for Artificial Intelligence, December 2025.

The AI Agent Identity Crisis

Current State of Digital Identity

Every identity system in production today, OAuth 2.0, SAML, OpenID Connect, Active Directory, was designed for one actor: a human logging into a service. Authentication flows assume real-time human decision-making. Credentials are stored on devices humans carry. APIs are optimized for UIs humans look at. Sessions are sized around human attention. And nothing in the stack verifies where a request is actually coming from.

Now introduce an AI agent. It runs continuously. It makes decisions across multiple services in the same second. It executes financial transactions. It operates at a scale humans can't match. And for security to mean anything, it needs to do all of this while acting independently, with authority that's scoped, revocable, and provable back to the human who delegated it.

Implementation Constraints

Financial Services: Current systems lack frameworks for secure AI trading bot deployment
Healthcare: Medical AI assistants face challenges accessing patient records within existing security protocols
E-commerce: Shopping automation requires frequent human approval, limiting efficiency gains
Enterprise: IT automation tools face credential exposure risks in current infrastructure

Across all four sectors, the same underlying gap appears in different costumes: authentication infrastructure assumes a human on the other end. Financial services can't issue scoped credentials to a trading bot. Healthcare can't prove a medical AI agent is acting within delegated authority at the record level. E-commerce can't safely decouple an agent's purchasing authority from a human credential. Enterprise IT can't avoid storing passwords the agent must read. These are not four separate problems requiring four separate solutions, they are one problem (agents don't have first-class identity) manifesting at four layers of the stack. The Kachyng identity layer addresses the common root.

IDX + KYA Architecture

The identity layer is built from seven interconnected components. The runtime control surface — Trusted Gateway, Zero-Credential Login, KYA runtime trust, and Rapid Revocation — is expanded in the Security and Revocation Model below.

1. **Agent Identity Framework**, Formal system for a human to create and operate multiple AI agents, each with an independent digital identity cryptographically bound to that human controller through Delegation Grant Tokens (DGTs). Includes ML-based trust scoring that adjusts each agent's autonomy over time.
2. **Trusted Gateway Architecture,** All agent interactions must originate from pre-registered IP ranges on Kachyng's secure infrastructure. Compromised tokens cannot be used from unauthorized networks, eliminating the largest class of credential theft attacks.
3. **Intelligent API Translation (Graphify**), Automatic conversion of legacy REST APIs to AI-optimized GraphQL responses. Eliminates UI metadata irrelevant to agents, reduces round-trips, and supports service-cooperative manifest discovery via standardized agent-api.json files.
4. **Payment Scope in Delegation Grant Tokens,** Per-agent payment authority expressed as scoped claims inside each DGT: spend caps, merchant allowlists, time windows, and per-transaction ceilings bound to the agent's identity rather than to a stored card. Biometric human oversight escalates on sensitive actions, and trust-based autonomy increases authority as agent reliability is demonstrated by KYA. Actual processor routing, MID portability, and rail selection are handled by PRX (Processor Exchange) and covered in its companion whitepaper.
5. **Zero-Credential Login Agent,** Specialized agents perform authentication on behalf of humans using credentials stored in encrypted vaults accessible only from trusted infrastructure. Credentials never touch user devices, eliminating phishing, keylogger, and malware attack surfaces entirely.
6. **Rapid Revocation Network,** Publish-subscribe architecture for global invalidation of agent authority. Supports per-agent, per-class, or system-wide revocation with cryptographic integrity and complete audit trails.
7. **Unified Adaptive Security,** All components controlled through biometric-secured mobile app with ML-based risk assessment. Behavioral pattern recognition detects anomalies; dynamic policies adjust trust scores and permissions based on demonstrated reliability.

How It Works: Procurement-Agent Walkthrough

Consider a mid-sized enterprise that deploys a procurement agent to replenish SaaS seats and cloud capacity as teams grow. The CFO needs automation without handing an AI agent a shared corporate credit card or a service-account password. Here is what happens across the IDX + KYA identity layer at each step.

Enrollment. The CFO authenticates once, then creates the agent as a distinct entity under the Finance team's organizational node. The agent receives its own cryptographic identity, not a reused human credential, bound to the CFO as controller through a Delegation Grant Token. The DGT encodes the org → Finance → CFO → agent chain, a $50,000 monthly spend cap, a merchant

allowlist (AWS, Snowflake, Okta, a pre-approved vendor list), and a 90-day validity window. (Agent Identity Framework; Hierarchical Delegation via DGT scope inheritance.)

Runtime isolation. The agent does not execute on a laptop or an arbitrary cloud VM. It runs from pre-registered IP ranges on Kachyng's secure infrastructure. If the agent's DGT were somehow exfiltrated, it could not be replayed from an attacker's network. The Trusted Gateway rejects any request not originating from Kachyng infrastructure. (Trusted Gateway Architecture.)

Vendor access. The agent needs to check seat counts on Okta and capacity on AWS. Instead of calling UI-oriented endpoints meant for human dashboards, Kachyng's Intelligent API Translation surfaces these as agent-native manifests via each vendor's agent-api.json, returning structured data without the UI cruft. Round-trips drop, and hallucination surface drops with them. (Intelligent API Translation / Graphify.)

Authentication to vendors. When the agent must sign in to a vendor portal that lacks agent-native auth, it does not carry a password. A Zero-Credential Login Agent performs the authentication from trusted infrastructure using credentials stored in an encrypted vault the agent itself cannot read, eliminating phishing and credential-theft surface entirely. (Zero-Credential Login Agent.)

Purchase. The agent identifies that the Engineering team needs twenty additional Snowflake seats at $2,400/month. The purchase request carries the agent's DGT, with payment scope intact: $2,400 is within the monthly cap, Snowflake is on the merchant allowlist, and the purchase is within the time window. The request is signed, verifiable, and traceable to the CFO's original delegation, not to a shared card. (Payment Scope in DGTs.)

Runtime trust evaluation. Before the purchase settles, KYA evaluates the action in real time against the agent's behavioral baseline. Purchase size, frequency, vendor, and timing all match the agent's pattern over the last 60 days. Trust score stays green, the action proceeds. Had the agent suddenly attempted a purchase from a merchant outside its allowlist, or at 3 a.m. when it has never transacted overnight, KYA would step up to biometric escalation requiring the CFO's approval on her phone. (Unified Adaptive Security / KYA runtime trust.)

The emergency case. A month later, the CFO sees an unfamiliar line item and wants to investigate. She taps "revoke" in the Kachyng console. Within seconds, every service that honors Kachyng-issued DGTs receives the revocation via publish-subscribe. The agent's next request, to any vendor, is rejected at the gateway. Auditable, fast, and not dependent on each vendor's individual credential-management quality. (Rapid Revocation Network.)

The full delegation lifecycle, creation, scope, execution, trust evaluation, and revocation, lives inside the identity layer, with no part of it depending on the vendor's native authentication stack. That is the IDX + KYA contract: a human delegates once, the infrastructure enforces the rest.

Security and Revocation Model

The walkthrough introduces four runtime controls one at a time. Consolidated, they are the control surface of the identity layer — four independent mechanisms that together make agent authority enforceable, observable, and revocable without depending on the downstream service.

Trusted Gateway: origin enforcement

Every request carrying a Kachyng-issued DGT must originate from a pre-registered IP range on Kachyng's infrastructure. A DGT exfiltrated to an attacker's network is useless: the gateway rejects it on origin alone, before any payload or policy check runs. This eliminates the largest class of token-replay and credential-theft attacks at the edge rather than in the application.

Zero-Credential Login: eliminated stored secrets

Where an agent must authenticate to a vendor portal that does not support agent-native auth, the agent never sees or carries the credential. A specialized Zero-Credential Login Agent performs the authentication from trusted infrastructure against an encrypted vault the agent itself cannot read. Phishing, keylogger, and malware attack surfaces that depend on a credential sitting on the agent's runtime are eliminated by construction.

KYA Runtime Trust: behavioral evaluation at action time

KYA evaluates each action against a behavioral baseline before it settles: purchase size, frequency, vendor, timing, delegation chain. In-pattern actions proceed. Out-of-pattern actions — a merchant outside the allowlist, a transaction at a novel hour, a scope extension the controller has not previously authorized — trigger step-up to biometric approval on the controller's device. Trust scores adjust dynamically over time, expanding or contracting the agent's autonomy without regrant.

Rapid Revocation: seconds-to-global invalidation

Revocation is publish-subscribe, not poll-per-service. When a controller revokes an agent, every Kachyng-integrated service receives the invalidation within seconds. Supported granularity: per-agent, per-class (e.g., every procurement agent under a given CFO), or system-wide. Every revocation event is signed and written to a tamper-evident audit log alongside the original delegation, so the before/after state of the agent's authority is reproducible for auditors, regulators, and incident responders.

These four mechanisms compose. A compromised token is rejected at the gateway, a stolen credential is unusable because it never left the vault, an out-of-pattern action is caught at runtime, and a revoked agent is globally dark in seconds. No single control is load-bearing alone; together they are the defense-in-depth the identity layer owes an enterprise deploying agents at scale.

What Is Live Today

The identity layer is in production. The capabilities below are demonstrable today under NDA against the live environment. Detailed walkthroughs, architecture diagrams, and transaction traces are available on request.

- Delegated agent identity issuance: human controller enrolls once, creates independent agents each bound by signed DGT.
- Scoped DGT-based authorization: spend caps, merchant allowlists, time windows, and hierarchical scope inheritance enforced inside the signed token.
- Trusted Gateway origin enforcement: requests rejected at the edge when not originating from pre-registered Kachyng infrastructure.
- Runtime trust evaluation: KYA scores each action against the agent's behavioral baseline, with step-up to biometric approval on anomaly.
- Revocation and audit evidence: per-agent and system-wide revocation propagated via publish-subscribe, with tamper-evident audit logs covering delegation, action, and revocation.

Compliance & Security Posture

The IDX + KYA architecture produces tamper-evident audit logs with cryptographic integrity by default, and is designed to carry compliance attestations for SOX, PCI DSS, HIPAA, GDPR, ISO 27001, SOC 2 Type II, and FedRAMP without bolt-ons. The deployment model supports multi-cloud (AWS, Azure, Google Cloud) with auto-scaling and edge presence, matching how regulated enterprises actually deploy identity infrastructure.

Intellectual Property & Patent Strategy

Portfolio Overview

Kachyng's identity architecture is protected by a patent family of 31+ assets with priority dates spanning 2010 to 2026, authored by the same inventor, Resh Wallaja, across two decades of identity work. The family covers the end-to-end autonomous commerce rail: mutual authentication,

identity aggregation, mobile-secured transactions, scope delegation, out-of-band approval, kill-switch controls, federated agent mesh, and ephemeral incident-scoped agents.

- Composition: 3 issued U.S. patents (priority from 2010), 18 provisionals filed June 2025–April 2026, 8 drafted and ready to file, 6 concept-stage. Utility conversion of the June 2025 cohort is scheduled for June 2026, establishing exclusivity through 2045 and beyond. Applicant: Kachyng, Inc. Inventor: Resh Wallaja (sole inventor throughout). Detailed filing information available under NDA.

Conclusion

The identity infrastructure built for humans does not carry over to autonomous agents. OAuth, SAML, Active Directory, and service-account patterns assume a human on the other end, a human attention span, and credentials humans can be trained not to share. An AI agent is a different actor: always on, multi-service, transacting in parallel, and delegating to sub-agents inside its own trust boundary. Treating it as a service account or a scoped human is the architectural mistake this paper is written against.

The IDX + KYA layer is built around that distinction. A human enrolls once. Each agent gets its own cryptographic identity bound to that human through a Delegation Grant Token. The DGT carries the scope — what the agent may do, where, for how long, and against what spend envelope. Every action runs from trusted infrastructure, is evaluated in real time against a behavioral baseline, and can be revoked globally in seconds. The full delegation lifecycle — creation, scope, execution, trust evaluation, revocation — lives inside the identity layer, with no part of it depending on the downstream service's native authentication stack.

This is Part 1 of 8. AGX (orchestration across enterprise systems), PRX (processor-agnostic settlement), and the broader Agentic Commerce Infrastructure thesis are covered in companion papers. Proof materials — production walkthroughs, architecture diagrams, filed patent claims, customer and processor evidence — are available under NDA.

Appendix: Competitive Landscape

The agent identity layer spans two overlapping markets: traditional Identity and Access Management (IAM) incumbents extending into non-human and agent identity (Okta, Microsoft Entra, Saviynt, CyberArk/Palo Alto), and a generation of purpose-built startups addressing agent-native identity, delegation, and verification (Aembit, Strata Maverics, Oasis, AuthZed, Vouched, Kite, t54 Labs). No single vendor today offers the combined IDX + KYA stack Kachyng provides: enrollment-time agent identity with signed delegation (DGTs), plus runtime behavioral trust and risk scoring bound to the same identity. Payment-rail players (Stripe, PayPal, Adyen, Visa TAP, Mastercard Agent Pay, Skyfire, Nekuda, Basis Theory) and fraud-detection vendors (Forter, Signifyd) are addressed in the PRX whitepaper; this appendix focuses on the identity layer.

The table below maps the most relevant players across these categories — their positioning, where their work overlaps with the IDX + KYA identity layer, and the architectural gap that distinguishes them from the integrated stack. Overlap reflects substantive alignment with the layer's core capabilities: signed agent identity, Delegation Grant Tokens, revocation, and runtime behavioral trust.

Company	Category / Positioning	Relevant Overlap	Distinguishing Gap
Okta	IAM incumbent extending into non-human and agent identity via Okta Identity Security posture and Auth0 agent protocols.	Agent identity guidance published 2025–2026; partnerships with AI platform vendors on delegated auth patterns.	Agent-native capabilities layered as bolt-ons over a human-centric IAM core; no signed delegation token, no runtime behavioral trust.
Microsoft (Entra)	Hyperscaler embedding agent identity into Entra ID, Copilot, and Azure AI Foundry; deep enterprise ownership.	Entra Agent ID previewed in 2025; Copilot Checkout + Mastercard Agent Pay integration announced November 2025.	Default enterprise identity plane, but architecturally tied to the Microsoft agent runtime; no cross-ecosystem delegation model or runtime behavioral trust.

Google (Cloud Identity, A2A)	Hyperscaler driving agent interoperability protocols (Agent2Agent, Universal Commerce Protocol) alongside Google Cloud IAM.	Launched A2A protocol 2025; joined Mastercard's UCP January 2026; pushing open agent discovery standards.	Protocol-layer focus; enforcement of identity, scope, and revocation at runtime is left to implementers.
CyberArk / Palo Alto Networks	Privileged access and secrets management incumbent; now part of Palo Alto's $25B acquisition (2025) positioning identity as platform security.	$25B Palo Alto acquisition announced 2025; Secure AI Agents product line launched for credential-based agent access.	Strong in secrets and privileged credentials, weaker in programmable agent delegation and scoped runtime authority.
Saviynt	Identity governance incumbent positioning SaviAI as unified governance across human, non-human, and AI identities.	$700M funding raise in 2025; launched SaviAI agentic governance workflows.	Governance-first framing; does not address real-time agent authorization at the moment of action.
Aembit	Workload and agent identity startup offering policy-based, short-lived access with OAuth / OIDC / SPIFFE / MCP support.	Programmable agent identity and runtime policy enforcement; integrations across AWS, Azure, GCP; Snowflake agentic workload partnership 2025.	Narrower scope than full IDX + KYA stack; no enrollment-time delegation chain binding agents to a human controller, no behavioral trust layer.
Strata Identity (Maverics)	Identity orchestration platform with purpose-built agentic identity flow: subject-actor OAuth binding, JIT provisioning, delegation chains.	Maverics Agentic Identity platform released 2026; subject-actor trust binding and delegation auditability.	Identity-first; no runtime behavioral trust layer and no payment-scope delegation.

Oasis Security	Non-human identity security platform covering AI agents across IaaS, SaaS, and on-prem with lifecycle governance.	Expanded agent coverage across ChatGPT, Copilot, Salesforce, GitHub throughout 2025; Fortune 500 deployments.	Discovery and governance focus; not positioned as an authorization runtime.
AuthZed (SpiceDB)	Fine-grained authorization specialist (Google Zanzibar-inspired) serving OpenAI and other AI-native customers.	Production-scale RAG authorization with OpenAI; explicit positioning against the "agent as me" pattern in 2025.	Authorization primitive, not a full identity stack; no issued agent identity, no behavioral trust, no payment scope.
Vouched	Identity verification startup repositioned around Know Your Agent (KYA): agent detection, credential verification, behavioral monitoring.	$17M Series A (September 2025) from Spring Rock Ventures; Agent Shield and Agent Bouncer product launches.	Inbound agent verification for websites; complementary to issuer-side identity rather than overlapping.
Kite (formerly Zettablock)	Agent Identity Resolution (AIR) with Agent Passport and blockchain-native stablecoin payment rails for autonomous agents.	$18M Series A (September 2025) from PayPal Ventures and General Catalyst; live Shopify and PayPal integrations.	Agent identity layer comparable in scope; thinner runtime trust scoring and delegation-grant architecture than IDX + KYA.
t54 Labs	Agentic finance "trust layer": agent identity verification, risk monitoring, and agent credit lines on blockchain.	$5M seed (February 2026) co-led by Anagram, PL Capital, Franklin Templeton; Ripple participation.	Institutional finance angle; narrower scope than a full agent identity platform.

Synthesis

The competitive picture is simpler than it looks.

The incumbents are structurally late. Okta, Microsoft Entra, CyberArk, Saviynt are all extending into agent identity from their existing access-management footprints, and every one of them has shipped an "AI agent" SKU in the last twelve months. They have distribution, which matters. What they don't have is architecture. Their platforms were built for static service accounts and human-delegated permissions. AI agents reason, delegate, and spawn sub-agents across trust domains. You cannot retrofit that. Microsoft's bundled Entra Agent ID is the most immediate incumbent comparator because it's free to existing Entra customers, but it remains architecturally tied to the Microsoft runtime and the same human-centric IAM core. The $25B Palo Alto–CyberArk deal and Saviynt's $700M raise tell you the incumbents know the window is short. They still can't close it in time.

The agent-native startups are each solving a slice. Aembit, Strata, Oasis, and AuthZed are going after enterprise IAM with the same buyer as the incumbents (the CISO) and a purpose-built architecture as the wedge. Kite and t54 ship an agent identity layer with thinner runtime-trust coverage. No current vendor unifies enrollment-time identity, signed delegation, runtime behavioral trust, and global revocation inside one identity layer.

The category is pre-standardization, which is the real "why now." NIST's AI Agent Standards Initiative and the MCP, A2A, AP2, TAP protocol competition are writing the rules in 2026. Whoever wins architectural influence over identity primitives captures durable value. Whoever sells the most seats does not. The enterprise deploying autonomous agents at scale is being forced to make a decision this year they'd rather defer. A new entrant wins by treating agent identity as a first-class primitive rather than a bolt-on, and by shipping while the incumbents are retrofitting.

That is what Kachyng is doing.

About Kachyng

Kachyng is building Agentic Commerce Infrastructure: the identity, orchestration, and settlement rails that let autonomous AI agents transact on behalf of humans and enterprises. Its patent-pending IDX + KYA identity layer gives every agent an independent cryptographic identity bound to a human controller — with scoped delegation, runtime behavioral trust, and global revocation in seconds. The architecture is protected by a patent family of 31+ assets with priority dates spanning 2010 to 2026.

Proof materials — production walkthroughs, architecture diagrams, filed patent claims, and customer and processor evidence — are available under NDA. Contact us to schedule a briefing.

Kachyng Inc. · 535 Mission Street, San Francisco

kachyng.com · info@kachyng.com