IDEMIA
PUBLIC
SECURITY

Identity in the Age of Agentics

Why Knowing Who is Behind the Agent is the Central Question for Establishing Trust in the Digital Economy

The internet was intended for people: for social connection, commercial transactions, to inform and educate, entertain us, and a host of other purposes. As the number of our interactions on the Web grow, we are developing an insatiable appetite for more and faster interactions. This is where agentic AI comes in. Agents are a continuation of our march toward automation online. As the Internet has grown, so have its non-human elements, with bots currently responsible for more than half of Internet traffic.[1] Meanwhile, AI deepfake and impersonation attacks, social engineering, and online fraudsters continue to exploit identity vulnerabilities *and* the predictability of our behaviors to rob us of what our identity entitles us to, and make us increasingly vulnerable to exploitation.

Agentic AI is building an Internet where *actions* matter more than presence—and where those actions may be taken without a human ever being actively involved. That shift fundamentally upends how trust works online.

Autonomous and semi autonomous AI agents are no longer experimental curiosities but are being actively used to enhance user experience. AI agents help us compare prices, move money, negotiate contracts, provision infrastructure, resolve customer claims, and orchestrate supply chains. Analysts now predict that by 2028, agentic AI will be embedded in a significant share of enterprise software, with agents making a meaningful proportion of business decisions autonomously.[2] Consumer adoption—especially in shopping and financial services—is rapidly following suit.

As we balance the excitement of the possibilities offered by agentic AI with its challenges, a basic question emerges:

Who, exactly, is behind the agent? And are they authorized to do what the person—or entity—has enabled them to do? In the age of agentics, identity is no longer a background security concern—it becomes the primary mechanism by which intent, accountability, and legitimacy are established.

In practical terms, the agentic world operates in two distinct modes:

- **Human present:** The individual is online and actively approving actions in real time. Identity, authentication, and consent are verified at the point of execution. The trust model resembles traditional digital commerce, but must still ensure strong identity verification and session integrity.
- **Human not present:** An AI agent acts autonomously under previously granted authority. The human is not actively participating at the point the transaction takes place. In this model, trust no longer comes from presence, it must come from ***cryptographically verifiable delegation.*** The agent must be provably bound to a verified human (or legal entity), and its permissions must be scoped, time-bound, and revocable.

As agentic systems scale, the second mode is becoming dominant, and identity assurances must evolve accordingly.

From Predictable Humans to Unpredictable Agents

For decades, digital trust has been built around human verification and human-centric patterns. People would log in, mostly during business or waking hours in their timezone; they hesitate and/or abandon processes in flight; they make mistakes. In short, even fraud has a human rhythm.

But AI agents break that assumption entirely.

A well-designed agent should act continuously, across time zones, at machine speed, without fatigue or direct supervision, and execute authorized tasks repeatedly and without error. By human standards, its behavior may look anomalous, and yet be perfectly legitimate. Conversely, a malicious agent can be engineered to look "normal" while doing extraordinary harm. As agentics grow, proving that there are verified individuals (or entities) behind them is critical to security and compliance.

This creates a dangerous ambiguity when it comes to fraud detection. Is an unusual transaction or set of behaviors inherently a sign of fraud, or simply automation working as designed? Is an action an expression of user intent, or the byproduct of misaligned autonomy and agentic execution? Without strong identity foundations, we cannot reliably tell the difference. Traditional security models attempt to detect abnormal behavior, but in an agentic economy, machine speed automation looks abnormal by human standards. The solution is not better anomaly detection alone, but verifiable identity and verifiable authority at execution time.

Instead of asking, *"Does this behavior look suspicious?"* systems must ask, *"Is this agent cryptographically authorized to perform this action on behalf of this verified individual or organization?"*

The challenge is no longer recognizing an agent just as an agent; it is affirming the true intent and authority of the human or organization behind it—every time an agent executes a task. This is where modernized digital identity solutioning such as verifiable digital credentials can play a central role in re-establishing trust and spotting potentially illicit, fraudulent, or simply mistaken activity.

[1] https://cpl.thalesgroup.com/about-us/newsroom/2025-imperva-bad-bot-report-ai-internet-traffic

[2] https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/articles/agentic-ai-enterprise-2028.html

Agentic Commerce and the Disappearance of Presence

Agentic systems enable commerce without human presence. Whether these include purchases, renewals, negotiations, arbitrage, or even dispute resolution, one can delegate these tasks entirely to software and agentic applications. The benefits may be obvious: speed, scale, personalization, and continuous participation in online markets. However, agentics also means that a critical component of trust has disappeared: *human presence as a trust signal*.

When a human clicks "buy," intent is assumed. When an agent does so, intent may be assumed, but indeed must be proven. Proving intent in a model where the human isn't actively engaged (or even present at all) in the transaction requires three critical elements:

- Proven identity of the human or entity behind the agent
- Explicit delegation defining what the agent may do, for what purpose, for how much, and for how long
- Cryptographic validation at the moment of execution

This is a structural shift in how trust is established online, as it is verified and validated on an ongoing basis—which inevitably brings us back to the question of identity. This framework will increasingly be applicable as agents themselves start creating their own subagents, where the "root" of the underlying authorized party can be affirmed down the agentic chain.

Governance on the Open Web, Delegation and Accountability

Inside enterprises, agents are supposed to operate within guardrails, often tied to corporate identity systems, governed by role-based access controls, licenses, policies, and audit requirements. These factors already have inherent risks, such as over-permissioned agents, orphaned agents (or orphaned identities), use of static credentials that are not updated, and user-related friction that companies are motivated to avoid.

On the open Web, standardized governance continues to be (re)defined and developed. An individual can deploy dozens of agents across shopping, investing, scheduling, negotiation, and content creation. Functionally, this means that the individual now resembles a small enterprise with multiple activities assumed to be authorized and automated. Yet the methods used for identity verification at the point of authorization continue to be fairly analog: using paper-based documentation, enabling a login, password, and potentially a weak second factor or email/text roundtrip.

At that point, systems struggle to answer basic questions. In effect, individuals now resemble small enterprises managing multiple digital actors. Enterprises have governance frameworks, role controls, auditing, revocation processes, but individuals currently do not. Agentic AI demands personal grade delegation frameworks that mirror enterprise governance: issuance of authority, scoping of permissions, revocation mechanisms, and traceable execution records.

Is the person behind the agent real? Are they the same person who authorized it? Is the agent permitted to perform *this specific action?* Is the counterparty interacting with a human, an agent, or a fabricated identity? Attackers thrive in precisely this kind of ambiguity. As the majority of fraud in online commerce and financial services has an identity nexus,[3] enabling comprehensive digital identity solutioning to track and trace agentic activities become equally critical.

We all know cases in an enterprise setting where unchecked delegation can be detrimental to security and have economic consequences. This is why companies are obligated to identify and verify beneficial owners of legal entities (the individuals who ultimately control or benefit from the entity's actions). Agentic systems now demand a similar concept for the individuals behind AI agents. Agents acting “on your behalf” are effectively exercising delegated authority, and must be affirmable as such.

If an agent can use delegated authority to act for me—spend money, commit resources, accept terms and deliver a service—then its authority must be **explicit, scoped, revocable, and traceable.** Doing so serves both security and legal priorities and requires reliable verification methods, ongoing monitoring, and clear delegation frameworks that reinforce accountability.

[3] **Impersonation fraud accounts for over 85% of all online fraud attempts,** making it the dominant fraud type and demonstrating that identity misuse is the central mechanism driving fraud across e-commerce and financial platforms; https://www.veriff.com/pr-news/veriff-identity-fraud-report-2026

A New Class of Identity Risk

The growth in agentic AI exacerbates existing risks and creates new vulnerabilities. Agents can not only potentially misrepresent their "owners," but also *themselves,* potentially being exploited by illicit actors, or other rogue/illicit agents serving different principals other than those authorizing them. Counterparties can misrepresent themselves to agents, leading users to be mistaken about who they are transacting with. Persistent agent memory can quietly accumulate sensitive identity data, creating high-value targets for attackers. Multi-agent systems introduce the possibility of forged delegation, where a sub-agent's request may never have been truly authorized.

Importantly, agents today often rely on static credentials, through API keys and cryptographic tokens that function as permanent "keys to the kingdom." These are precisely the kinds of credentials attackers steal. With continued use of analog identity capture and verification real-time verification of agents as they perform tasks remains cumbersome and ineffective—not to mention the rise of AI deepfake, synthetic identity, and social engineering attacks that make such credentials vulnerable to theft. Stolen or compromised credentials remain the most common initial access vector in real world breaches, according to Verizon's 2024 Data Breach Investigations Report[4]. In short, **authentication that only proves knowledge rather than identity itself does not scale to autonomous systems.**

[4] https://www.hipaajournal.com/verizon-2024-data-breach-investigations-report/

Verifiable Digital Credentials (VDCs): Modernized Identity for Web-Native Activities

Encouragingly, a scalable identity foundation has already emerged, as governments and trusted institutions (especially in highly regulated industries) are testing verifiable digital credentials: encrypted, user-controlled, digital versions of passports, driver's licenses, professional qualifications, and other authoritative documents. These credentials are issued by trusted entities, stored under user control, and presented selectively.

This architecture makes VDCs uniquely suited for use in agentic systems. A human can issue narrowly scoped, reusable credentials to an agent without exposing raw personal data or payment instruments. The credential can encode the limits of authority itself, amount thresholds, permitted counterparties, duration of validity, while preserving privacy through selective disclosure.

VDCs are more secure, more private, and more reusable than traditional form factors or paper-based identity. As such, they are uniquely placed to address the needs of agentic systems. Adoption continues to accelerate worldwide, and global regulators and policy makers are increasingly looking to VDCs to modernize regulatory obligations such as Bank Secrecy and Know Your Customer (KYC) activities. Estonia and Australia have long-running national digital ID programs. In the United States, more than 20 states now support app based mobile driver's licenses accepted at TSA checkpoints and federal agencies—the majority of which are issued by IDEMIA Public Security on behalf of states. Europe is preparing to launch the EU Digital Identity Wallet, which every EU member state must offer by the end of 2026.[5]

[5] https://commission.europa.eu/topics/digital-economy-and-society/european-digital-identity_en

Reestablishing Trust: Digital Identity as the Control Plane for Agentics

Imagine authorizing an AI agent to purchase airfare on your behalf within a $500 limit over the next 30 days. Rather than sharing your card details or full identity record, you issue a verifiable digital credential that encodes:

- Your verified identity
- Your verified payment instrument
- The permitted action (purchase airfare)
- A monetary cap
- A defined validity period

When the agent executes the transaction, the merchant can cryptographically verify both your identity and the scope of delegated authority. If the action exceeds those limits, it fails automatically. If all elements are assured, payment instructions may be executed, and verified payment mechanisms are engaged. Every step is recorded and provable, creating accountability without requiring constant human oversight.

As AI agents proliferate, digital identity becomes the control foundation that makes autonomous agents trustworthy and secure. Agents can be affirmed and bound to real, verified humans and organizations. Digital Identity—and VDCs—can allow transparent delegation, ensure rules and revocation when needed, and proactively identify and disrupt fraud. When enabled, VDCs ensure the persistence of accountability when humans are no longer present at the point of action. Without it, agentic systems will continue to oscillate between two failure points: being so permissive that they invite abuse, or so constrained that their value is undermined.

Conclusion

The explosion in Web-based commerce and interactions has revealed that trust does not inherently emerge from scale. The growth of online engagement, where bots proliferate and the number of autonomous interactions grow, requires trust to be re-engineered on the open Web. AI agentics will continue to accelerate online traffic in almost every area in which we engage: speed, reach, economic impact, and risk.

Enabling identity as critical infrastructure on the Web ensures durability through the cryptographic expression of who is acting, with what authority, under whose delegation, and for whose benefit.

Trust in the agentic economy will not emerge from behavior monitoring alone. It will be built on **verifiable identity, transparent delegation, scoped authorization, and tamper-proof execution records.** Only then can autonomy scale without sacrificing accountability.

For more information, contact us at **info@ps-idemia.com**

IDEMIA PUBLIC SECURITY

NA.IDEMIA.COM