CHAINIT

EXECUTIVE WHITE PAPER

Agentic Compliance Begins With the Human Behind the Agent

Author
Jeremy K. Blackburn, Founder and Chief Executive Officer, ChainIT Inc.

When an AI agent initiates a payment, the money still belongs to a person or an organization. The parties to that transaction still have identities, restrictions and obligations. Introducing software into the relationship does not make those facts disappear.

An agent can present valid credentials and communicate through an approved platform while the business receiving its instruction cannot establish the real person or legal entity behind it. For a corporate agent, the same system may be unable to establish whether the human who delegated the task had authority from the business to do so.

That raises a question the agentic economy needs to answer:

If there is no verified human connection to the agent, who is compliance being conducted on?

In my previous article, I discussed why AI agents need a verifiable human connection to establish authority. Compliance is the next consequence of that relationship. We need to know who the agent represents, whether that person had authority to delegate, and whether the proposed activity is permitted for the parties involved.

For individuals, the connection must reach the person on whose behalf the agent acts. For organizations, it must reach the legal entity and the verified human whose authority supports the delegation, together with other owners or controllers the applicable rules require us to identify.

A Recognized Agent Can Still Conceal an Unknown Principal

Agent identity serves an essential security purpose. A business should be able to recognize the software making a request, validate its credentials and determine what systems it can access.

Customer compliance asks additional questions. A software identifier alone cannot establish the age of a purchaser, the identity of a customer, the ownership of a company or the sanctions status of the person directing a transaction.

Transaction monitoring and behavioral controls can still detect suspicious activity when identity information is incomplete. But when a rule depends on who the customer or another relevant party is, that missing identity becomes a substantive gap in the decision.

Bad actors have every reason to exploit that gap. An agent can become a professional-looking intermediary between a prohibited party and a legitimate business. If verification stops at the software, the person or organization behind it gains another place to hide.

A legitimate platform does not make every user legitimate. A properly authenticated agent can carry instructions from someone who should never have been permitted to initiate the transaction.

Verifying the human gives compliance a real subject to evaluate. It does not certify good intentions. A correctly identified person can still commit fraud, act for someone else or attempt a prohibited transaction. The verified connection allows the institution to evaluate those risks against an identifiable party and investigate the relationship when something is wrong.

The Verification Gap in Current Agent Frameworks

Current agent frameworks increasingly support user consent, delegated access and business sponsorship. Those capabilities help secure automation. The compliance gap appears when a deployment treats that technical relationship as sufficient evidence of the underlying person or business.

The Model Context Protocol, or MCP, defines OAuth-based access to protected tools and resources. Its authorization server can interact with a user and issue access tokens, but the server's implementation is outside the specification's scope. The identity evidence supporting that account must come from the deployment. [1]

The Agent Payments Protocol, or AP2, records user authorization through signed mandates and binds approvals to checkout and payment activity. Its current specification requires deterministic validation. Its authorization framework also places credential issuance outside that document's scope. Compliance therefore depends on what the credential issuer or agent provider actually establishes about the principal and the authority behind the approval. [2,3]

The Universal Commerce Protocol, or UCP, supports OAuth account linking, user consent and external identity providers. Its B2B example also leaves customer eligibility and payment terms to the merchant to resolve at runtime. Microsoft Entra Agent ID assigns business sponsors to agent identities. These mechanisms provide account relationships, access controls and organizational accountability. The remaining question is what evidence establishes the real person, the represented legal entity and that person's authority to delegate the particular activity. [4,5]

An authenticated account may belong to a real, verified customer. A sponsor may have the required corporate authority. A credential may carry reliable evidence of both. The institution accepting the action must be able to establish those facts instead of assuming them from the presence of an account, sponsor or signature.

These frameworks can carry or reference stronger identity and authority evidence. Adopting a protocol alone does not create that evidence. Where the relationship has not been established, the system can validate the agent's request while leaving the individual or business behind it unverified for the applicable compliance obligation.

That is the gap: a valid technical request with no proven compliance connection to the accountable person or business.

Individual Compliance Must Reach the Individual

Consider a personal agent instructed to transfer money. Where customer due diligence applies, the provider needs a verified customer relationship and evidence that the agent is acting within that customer's actual authorization. An account name attached to an agent is only useful if the relationship behind it has been established.

For U.S. banks subject to 31 CFR 1020.220, the Customer Identification Program must use risk-based verification sufficient to establish a reasonable belief about the customer's true identity. The rule addresses the customer, whether an individual or a legal entity. Recognizing the software making a request does not by itself establish that identity. [6]

The same logic applies to an age-restricted purchase. The eligibility question concerns the person the restriction applies to. Running the instruction through software cannot change that person's age or transfer another person's eligibility to them.

Identity verification, authentication and delegation have different jobs. Verification establishes the person's identity to an appropriate level of assurance. Authentication checks the credentials or factors bound to that identity. Delegation records what the agent may do on that person's behalf.

Those relationships need protection against impersonation and account takeover. They also need limits, expiration and revocation so an agent cannot continue relying on authority the person has withdrawn.

Routine activity can proceed automatically within a valid mandate. Renewed human involvement should occur when the mandate, risk or governing policy requires it. The human connection provides the foundation for responsible automation.

Corporate Compliance Requires More Than a Verified Employee

Organizational compliance must establish several distinct relationships.

A company may be legally registered. An employee may have a verified identity and a working corporate account. Those facts still leave an unanswered question about whether that employee can authorize an agent to spend company money, bind the company to an agreement or disclose protected information.

The authority must come from the organization and cover the activity being delegated. A procurement manager's purchasing authority, for example, may be limited by amount, supplier, business purpose or a required second approval. Evidence may include governing documents, resolutions, authorized-signatory records or bounded delegations, depending on the organization and transaction. The agent's authority must remain within those boundaries.

Consider a purchasing manager with a $10,000 approval limit who gives an agent a $40,000 mandate. The person can be correctly identified, the corporate account valid and the recorded instruction authentic. The instruction still exceeds that manager's delegated authority. The missing evidence is the additional corporate authorization needed for that amount.

The authorized representative is also a different compliance subject from the company's beneficial owners and controllers. Verifying an employee does not resolve the ownership questions that may apply to the business. Nor does verifying the company resolve every question about the individual acting for it.

FATF's Interpretive Note to Recommendation 10, paragraph 4, calls for financial institutions to establish both the identity and the authority of a person acting for a customer. Paragraph 5 separately addresses legal entities and their ownership and control. FATF sets international standards implemented through national frameworks. Applying those principles to agent delegation is the architectural argument here. [7]

For covered U.S. financial institutions, 31 CFR 1010.230 generally requires identifying and verifying each individual, if any, who directly or indirectly owns at least 25 percent of a legal entity customer's equity, together with one individual having significant responsibility to control, manage or direct the entity, subject to exclusions and exemptions. FinCEN permits reliance on customer-provided information under specified conditions. The agent's delegator may be neither an owner nor that control person. [8]

OFAC provides a concrete example. FAQ 400 explains that transactions involving a blocked person are generally prohibited absent OFAC authorization, even when that person represents a company that is not blocked. For a transaction subject to U.S. sanctions, screening the company alone can therefore miss a prohibited human participant. [9]

Ownership is a separate inquiry. OFAC's FAQ 398 explains that control alone does not automatically block an entity under the 50 Percent Rule. The representative's involvement can still matter even when that ownership test does not block the company. [10]

An organizational agent needs a verifiable chain from the company to its authorized human representative and onward to the agent. The compliance assessment must also reach the other relevant parties and facts required for the activity. A corporate email address, a directory role or a business sponsor assignment is useful only to the extent that the supporting evidence establishes the relationship the decision requires.

Compliance Has to Hold When the Agent Acts

Even a well-established relationship can become outdated.

Imagine a treasury manager authorizing an agent to pay an approved supplier. Before the payment is released, the supplier's eligibility changes, the company revokes the delegation or the payment destination is altered.

The agent may still have an unexpired credential. The earlier verification may have been entirely correct. The proposed payment now presents a different decision.

Agentic Compliance should evaluate the relevant identity, authority and compliance conditions at the point where the action can still be stopped. For a payment, that means before release. For an agreement, it means before the organization becomes bound. For protected data, it means before disclosure. This is an execution design principle for making applicable obligations effective in an automated workflow.

That decision must cover the actual instruction. An approval for one payee and amount cannot simply follow an agent into a changed transaction. If material details change, the applicable checks and approvals must be reconsidered.

Existing verification evidence can be reused under appropriate risk and freshness rules. A new biometric interaction is not inherently required for every action. FinCEN's February 2026 relief permits covered institutions to forgo repeated beneficial-owner identification and verification solely because an existing legal entity customer opens another account, while retaining initial checks, reliability triggers and risk-based ongoing diligence. That relief concerns account openings; it illustrates evidence reuse rather than prescribing agent checks. [11]

The architecture should establish which evidence remains reliable and which conditions need a current check. Where required evidence is unavailable, the execution system should hold or deny the action according to the applicable policy. Missing information should never quietly become permission.

The Agent Must Operate Within Enforceable Rules

The agent can help assemble information, assess options and propose a course of action. The system that releases money or grants access needs a separate, enforceable decision.

A model's conclusion that a transaction appears acceptable cannot override a revoked delegation, a prohibited party or a mandatory approval. Those controls must remain effective even when the agent is mistaken, manipulated or compromised.

Human approval also has a defined role. An authorized person may approve an exception permitted by company policy. Human approval remains subject to legal prohibitions and any applicable requirement for regulatory authorization.

Oversight depends on preserving the relationship between the evidence, the rule and the outcome. A regulator or internal reviewer should be able to establish which principal was represented, how authority reached the agent, what conditions were evaluated and what the execution system actually did.

A detailed log of an unidentified agent's activity can explain the sequence of events. Effective accountability also requires evidence connecting those events to the relevant people and organizations. Technology supports the regulated institution's compliance program; it does not transfer the institution's obligations to the agent.

Privacy and Accountability Can Work Together

A verifiable human connection can be designed with selective disclosure. The agent does not need a copy of the person's identity documents or biometric information to perform every task. A merchant checking age eligibility may need a reliable eligibility result rather than a full identity dossier.

The necessary information depends on the role and the obligation. A regulated institution may need underlying identity and ownership records that an ordinary merchant does not. The design

should give each party the evidence it needs, protect the underlying data and preserve lawful access for oversight.

The useful distinction is between limiting unnecessary disclosure and leaving the accountable party unknown. We can protect personal information while preserving a verifiable connection to the party whose eligibility matters.

Where ChainIT Makes the Connection

At ChainIT, our approach begins with the human and the organization behind the instruction. We connect Agentic Identity, Agentic Authority, Agentic Compliance and Agentic Commerce through evidence that remains relevant to execution.

Agentic Identity establishes who is represented. Agentic Authority establishes what the agent may do for that principal. Agentic Compliance evaluates the relevant parties and proposed activity against applicable requirements. Agentic Commerce carries those relationships and decisions into transactions. Each depends on evidence that the next stage can use.

ChainIT ID connects the individual to verified identity evidence. Our five-factor authentication approach, including biometrics and liveness checks, helps protect that connection. ChainIT Org ID establishes the organizational context. The authority relationship then determines whether the identified individual can delegate the proposed activity for that organization.

Our Validated Data Tokens carry verified attributes and supporting evidence. The decision service evaluates the specific instruction against the applicable authority and compliance conditions and returns allow, hold, require approval or deny. The payment or execution system enforces that result before the action proceeds.

The evidence must connect the principal, delegation and decision to the actual instruction and its eventual outcome. Any required legal response must also be carried through in the institution's workflow. A record of approval that cannot be matched to the executed payment leaves the central question unresolved.

This approach works alongside customer due diligence, sanctions screening, transaction monitoring and the institution's other controls. It can support the evidence consumed by agent protocols and enterprise identity systems. The applicable rules and transaction facts determine the compliance result.

The gap ChainIT addresses is the verifiable relationship between an agent and the real individual or business it represents, including the authorized human delegation, carried into an enforceable decision at execution.

Automation Must Preserve the Accountable Party

The commercial value of agents comes from their ability to act for us. That value depends on businesses being able to determine whose activity they are accepting and whether it is permitted.

For an individual, the compliance record must reach the actual person represented. For a company, it must establish the entity and the authorized human delegation, with ownership and control information where required. It must then apply the relevant rules to the transaction being executed.

Without that connection, we give bad actors another place to hide and compliance teams a record that stops short of the party they need to understand.

We should build agentic commerce so the human can step away from routine execution while the evidence of identity, authority and compliance remains available and enforceable.

An agent can act on someone's behalf. Compliance must still reach that someone.

Regulatory and Technical Sources

Sources checked October 31 2026. The discussion of agent architecture is the author's application of these authorities and specifications. Obligations depend on the institution, activity and jurisdiction; the cited materials do not prescribe ChainIT's particular implementation.

1. Model Context Protocol Authorization. Specification dated July 28 2026. See Purpose and Scope and Roles for protected-resource access and the boundary around authorization-server implementation.

2. AP2 Agent Authorization Framework. See Mandate Delegation and User Credential. The framework addresses user consent and authorization; credential issuance is outside that document's scope.

3. AP2 Agent Payments Protocol. Version 0.2. See Agentic vs Non-Agentic, Mandates and Verification for transaction binding and deterministic validation.

4. Universal Commerce Protocol Identity Linking. Schema dated August 25 2026. Describes OAuth account linking, user consent and external identity-provider connections.

5. Microsoft Entra Agent ID Administrative Relationships. Describes owners, sponsors and managers. Business sponsorship provides lifecycle accountability; the article distinguishes it from evidence of authority for a particular economic act.

6. 31 CFR 1020.220 Customer Identification Program Requirements for Banks. See paragraph (a)(2) for risk-based customer identification and verification. Additional verification of entity representatives under (a)(2)(ii)(C) is conditional, not a universal rule for every employee.

7. FATF Recommendations. Updated June 2026. Interpretive Note to Recommendation 10, paragraphs 4 and 5, printed pages 66 to 68. FATF sets international standards implemented through national frameworks.

8. FinCEN Customer Due Diligence Rule FAQs. Consolidated guidance updated May 6 2026 on 31 CFR 1010.230. See sections A and B for scope, ownership and control, verification, reliance and ongoing diligence.

9. OFAC FAQ 400. Addresses transactions involving a blocked individual representing a non-blocked entity, including the OFAC authorization qualification.

10. OFAC FAQ 398. Distinguishes ownership under the 50 Percent Rule from control and from a blocked representative's participation.

11. FinCEN Order FIN-2026-R001. February 13 2026. Exceptive relief from repeated beneficial-owner verification at subsequent account openings, retaining initial identification, reliability triggers and risk-based ongoing diligence.