AI Agents Need a Verifiable Human Connection

Identity Authority and Compliance for Agentic Commerce

By Jeremy Blackburn, Founder and CEO, ChainIT Inc.

Revised edition | 31 October 2026

chainit.com

Hand a valet your car keys, and the task is understood. Park the car. Bring it back when asked. The keys give him access to the vehicle. The arrangement defines what he may do with it.

Nobody confuses possession of the keys with ownership of the car or unlimited authority over it.

Now give an AI agent a login, a wallet or an API credential. It can reach an account, retrieve information and perhaps move money. But what travels with that access to prove whose interests it represents, who authorized it and where its authority ends?

That question sits underneath the permission problem Karen Webster describes in her PYMNTS article, “Why AI Agents Will Live or Die on Permission.” She makes the commercial case for businesses controlling how agents interact with their systems and who bears the consequences when something goes wrong.[1]

I agree. I would start one step earlier.

Before a business can rely on permission, it needs to verify its source. The person or business granting permission must be identifiable. Whoever delegates authority must actually have the right to delegate it.

That verifiable human connection is the missing piece ChainIT provides: the link between an agent and the individual, or the legitimate organizational authority, behind its actions.

Agentic Authority is a verifiable delegation permitting an AI agent to act for an identified individual or organization within defined limits. Establishing it requires evidence of the principal’s identity, the right to delegate, and the scope and current validity of the grant.

The human connection is the foundation of a sequence that agentic commerce has to get right: identity, then authority, then compliance.

Knowing the Agent Is Only Half the Job

Knowing which agent is making a request is useful. So is knowing who built it, which credentials it holds and whether its software is operating in an approved environment. Those answers help establish what the system is dealing with.

The commercial question goes further. On whose behalf is it acting?

Agentic Identity, as I use the term here, connects the identified agent to the verified individual or organization it represents. The company that developed an agent may have no authority over the customer's money. The employee who installed it may have no authority to sign the company's contracts.

We should recognize the work already underway. Google's Agent Payments Protocol uses signed mandates to capture purchase instructions and approvals. Microsoft's Entra Agent ID assigns business sponsors to agents. These are meaningful steps toward connecting agents to accountable people and instructions.[2,3]

A credential can be valid even when the real-world authority behind it is unproven. A named sponsor can be genuine and still lack permission to delegate a particular transaction.

What must be verifiable is the relationship between the agent, the individual or legal entity it represents, and the source, scope and current validity of the authority being exercised. ChainIT connects those facts to the specific action and the controls governing its release.

The connection becomes commercially useful when the institution receiving the request can verify it for itself.



A Personal Agent Needs More Than a Login

Suppose I ask a personal AI agent to find a flight to Chicago. It finds one, purchases a nonrefundable ticket, books a hotel and buys travel insurance.

It may have planned a perfectly sensible trip. It may also have spent money I never authorized it to spend.

Now suppose I explicitly authorize a particular flight, up to $1,200, using a designated payment method. The delegation has substance. It identifies an action, a limit and a source of funds. The airline and payment provider still need a reliable way to connect that instruction to me and to the purchase actually submitted.

The same distinction matters when an agent reads financial records, shares personal information or changes account settings. Permission to inspect a balance says little about permission to transfer it. Permission to book an appointment does not settle which records may be disclosed along the way.

A useful human connection therefore has to preserve what I authorized, for which purpose, within which boundaries, and for how long. It also has to let me withdraw that authority.

That does not require me to sit beside the agent approving every click. I can delegate routine work in advance. Explicit, enforceable limits let the software operate within that delegation while I do something else.



A Business Agent Has Another Relationship to Prove

For a business, verifying the human is only the beginning.

Imagine a procurement manager authorized to buy up to $10,000 of approved supplies for a project. She gives an agent a task, and it finds a $50,000 deal that looks too good to miss.

Her identity can be authentic. Her employment can be real. Her instructions can be accurately recorded. The proposed purchase still exceeds the authority she was given.

A company email address has never been a blank check.

The business may also restrict which suppliers she can use, which contracts she can sign and whether she can delegate any of that authority to software. The agent's mandate has to fit inside those restrictions. Dividing the purchase among several agents cannot legitimately turn the project's $10,000 limit into five separate budgets.

The connection must therefore reach beyond the employee to the organization and its legitimate governance. The organization can be the principal. The people, roles, resolutions and policies through which it grants authority must be verifiable.

That relationship changes. An employee leaves. A purchasing limit is reduced. A contract expires. A delegation is revoked. A credential that worked yesterday may still authenticate today, even though the authority needed for the proposed action has ended.

The practical question is whether the agent has current authority when the business is about to incur the obligation.

Agentic Compliance Depends on the Relationship and the Action



Agentic Compliance evaluates whether a proposed action satisfies the applicable requirements, using the relevant identity, authority and transaction evidence.

A verified identity does not create authority. Valid authority does not make every transaction compliant.

An authorized payment can still require a hold because a destination changed, relevant screening information is stale or an additional approval is required. Those conditions have to be evaluated when the action is about to occur.[1,2]

Without a full understanding of the underlying relationship, the evaluation is missing essential context. Whose money is being used? Which business is making the commitment? What limits govern the person or agent submitting it? What purpose was approved?

Agentic Compliance must evaluate the relevant people and entities as well as the proposed action. Depending on the applicable requirements, those subjects can include the customer, an authorized representative, beneficial owners or control persons, and the recipient or counterparty. Validating the software's credentials cannot substitute for establishing the identities and relationships those checks require.[4,5,6,8]

The decision must also affect execution. If the agent can bypass a failed check by using another credential or route, the restriction has little practical value.

A useful control is enforced where the money, data or commitment is released. It must cover the execution paths available to the agent within that workflow.

Existing Obligations Follow the Principal

The human connection also has an existing regulatory foundation. Delegating a task to software does not, by itself, discharge the relevant parties' obligations.

FATF's international standards call for financial institutions to verify that a person acting on a customer's behalf is authorized, and to identify and verify that person. These are standards for national implementation, rather than a single law governing every agent.[4]

In the United States, bank Customer Identification Program requirements address risk-based verification of customer identity. Separate customer due diligence rules require covered financial institutions to identify and verify specified beneficial owners of legal entity customers, subject to exclusions and relief.[5,6]

FinCEN's February 2026 order reduces repeated beneficial-owner verification at additional account openings while preserving initial requirements and risk-based updating. The practical requirement is to have the evidence the applicable rules require when it is needed.[7]

The customer, an authorized representative, and an ownership or control person are distinct roles. Identifying one does not automatically establish the others or prove authority for a particular action.

OFAC's guidance separately encourages risk-based sanctions controls. Valid identity and delegation do not make a prohibited transaction permissible.[8]

For agentic workflows, the design task is to preserve these distinctions and carry relevant evidence into the controls governing the proposed action. Cryptographic binding can support that task; these sources do not prescribe it as a universal requirement.

Where ChainIT Connects the Pieces

At ChainIT, we connect verified identity to authority, applicable controls and evidence of the resulting action. The agent's request can then be evaluated against a verifiable relationship to the person or business it represents.

ChainIT ID establishes the individual through verification against authoritative identity records, together with biometric matching and liveness controls. ChainIT Org ID establishes the business. Together, they identify the person and organization whose relationship and authority must then be verified.[9]

ChainIT records an organization's grants of authority and their constraints in an Authority Resolution Pactvera, or ARP, a corporate resolution executed on Pactvera, our contracting platform for identity-verified agreements. It connects the person to the entity and provides a basis for checking whether that person may act for it. ChainIT verifies and operationalizes that authority; the underlying right comes from the organization's legitimate governance and applicable legal arrangements.[10]

Our Provable Authority architecture extends that chain to the agent and the exact action. It specifies delegated scope, limits, validity and revocation conditions, with independent rules governing whether execution may proceed. Where approval is required, it is bound to the actual transaction details. Under that design, changing the payee or amount invalidates the original transaction approval.[11]

Applied to our procurement example, the architecture requires additional authority for the $50,000 purchase. Authenticating the manager and the agent does not expand the manager's $10,000 spending authority.

Provable Compliance evaluates the relevant evidence and policy as a separate requirement. The resulting records preserve the basis of each decision and its connection to the action.[12]

AI can search, compare, recommend and coordinate. The controls governing whether its proposed action is allowed must remain outside the model's discretion. An agent should have no ability to expand its own mandate because doing so would make the task easier.

This is what we mean by the human connection: a relationship that can be verified, constrained and carried through execution. It gives a counterparty a basis for deciding whether to accept the agent's request.

The Proof Has to Travel With the Agent

Agentic Commerce is commercial activity in which AI agents carry out tasks and transactions for individuals or organizations, often across multiple systems. Consider a business agent that negotiates with a supplier's agent and then asks a payment provider to settle the invoice.

Each participant sees part of the transaction. The payment provider may recognize the credentials. The supplier may recognize the customer. The agent platform may hold the original instruction. The authority connecting them needs to remain verifiable through those handoffs, with the original limits intact.

In ChainIT's design, an agreement executed on Pactvera can settle through ChainIT Pay, allowing the authority checks and applicable limits to govern payment release. A business can also retain its existing payment provider when its payment workflow is integrated to enforce the approved instruction and its limits before funds are released, and to link the resulting payment evidence back to that authorization.[13]

The authorization I give to an agent does not give it unrestricted rights over someone else's systems. A bank, merchant or data owner still applies its own access conditions and controls. Both sides of the relationship matter.

Verification should also disclose only the information needed for the decision. A counterparty may need evidence that an authorized representative approved a purchase without needing that representative's entire identity file.

Afterward, the evidence should show which authority supported the action, what limits and approvals applied, the exact instruction submitted and what actually occurred. ChainIT preserves those linked records in an immutable, append-only history. Their value as evidence comes from verifiable provenance and the connection between authorization and execution, as well as protection against alteration.[11,12,13]

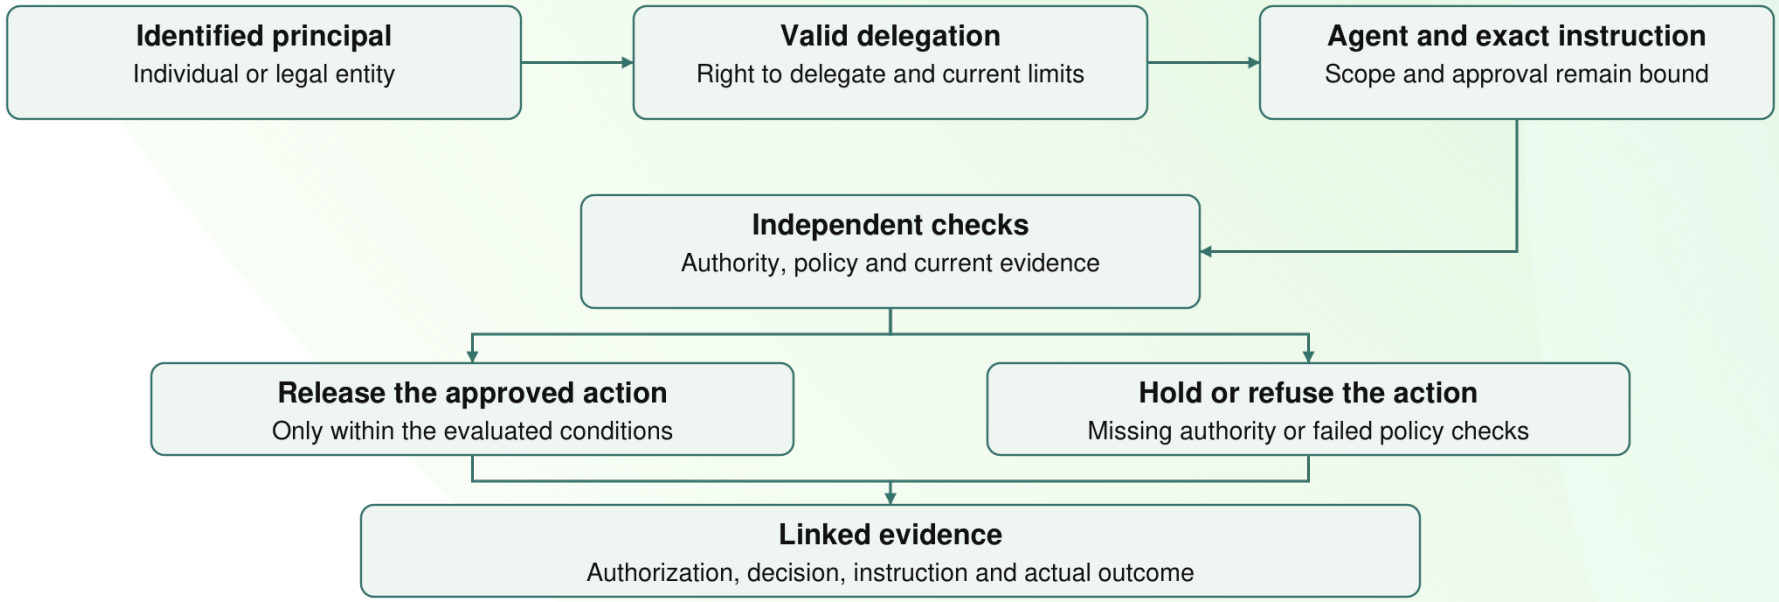
That is the basis for investigating a disputed transaction and distinguishing an authorized decision from an agent exceeding its mandate. The record does not, by itself, decide legal liability or relieve an agent provider of responsibility for its own failures. It gives the parties evidence with which to resolve those questions.

Authority Makes Automation Usable

When a business cannot verify the authority behind an agent's request, manual review and broad restrictions become rational responses. Both cost time. Both can prevent legitimate business from getting done.

A verifiable connection gives the receiving business a more useful choice. It can accept requests within defined authority, require additional approval when circumstances warrant it, and refuse requests that fall outside the mandate.

That is the commercial value of the infrastructure ChainIT provides. It gives institutions a basis for allowing useful automation while retaining control over the obligations they accept.



The decision is enforced before release and recorded with the resulting outcome.

The human does not have to be present at every step. The evidence of human or organizational authority does.

We already understand this when we hand someone our car keys. Agentic commerce needs to preserve the same distinction at machine speed when an agent moves money, releases data or creates an obligation across systems.

That is the human connection ChainIT brings to agentic commerce.

References

Primary sources for the article and its regulatory and product claims. Sources reviewed 31 October 2026.

1. Karen Webster. Why AI Agents Will Live or Die on Permission.
PYMNTS, 28 October 2026. The permission and commercial accountability argument that prompted this article.
Read the article

2. Google Cloud. Powering AI commerce with the new Agent Payments Protocol.
16 September 2025. AP2 uses signed mandates to capture purchase intent and approval.
AP2 announcement

3. Microsoft Learn. Administrative relationships in Microsoft Entra Agent ID.
Accessed 31 October 2026. Distinguishes agent owners, business sponsors and managers.
Agent ownership and sponsorship

4. Financial Action Task Force. The FATF Recommendations.
Amended June 2026. Interpretive Note to Recommendation 10, Section B, paragraph 4, addresses the authorization and identity of a person acting for a customer. These international standards are implemented through national measures.
Official recommendations

5. Electronic Code of Federal Regulations. 31 CFR 1020.220.
Customer Identification Program requirements for banks. Subsection (a)(2) requires risk-based customer identity verification. Subsection (a)(2)(ii)(C) addresses additional checks on persons with account authority or control when the bank cannot otherwise verify a nonindividual customer’s identity.
31 CFR 1020.220

6. Electronic Code of Federal Regulations. 31 CFR 1010.230.
Beneficial ownership requirements for legal entity customers. Applies to covered financial institutions and specified legal entity customer relationships, with exclusions and exemptions. Ownership and control are separate from authority for a particular transaction.
31 CFR 1010.230

7. Financial Crimes Enforcement Network. FIN-2026-R001.
Exceptive relief order, 13 February 2026. Permits covered institutions to limit beneficial-owner identification and verification to the first account opening, later reliability concerns, and circumstances required by risk-based ongoing due diligence.
Official FinCEN order

References Continued

8. Office of Foreign Assets Control. A Framework for OFAC Compliance Commitments.

2 May 2019. Describes a risk-based sanctions compliance program. Applicable sanctions requirements must be assessed separately from identity and delegation.

OFAC framework

9. ChainIT. Identity and organization verification.

Product materials, accessed 31 October 2026. Describes ChainIT's identity foundation and the roles of ChainIT ID and ChainIT Org ID.

ChainIT platform | ChainIT ID | ChainIT Org ID

10. ChainIT. Authority Resolution Pactvera.

Product materials, accessed 31 October 2026. Describes the relationship between organizational authority and identity-verified agreements on Pactvera.

Authority Resolution Pactvera

11. ChainIT. Provable Authority.

Architecture announcement, 2 September 2026. Describes verified identity, delegated authority and approval for the exact action.

Provable Authority announcement

12. ChainIT. Provable Compliance.

Architecture announcement, 8 September 2026. Describes evaluation of current compliance evidence and policy separately from authority, with records of the resulting decision.

Provable Compliance announcement

13. Jeremy Blackburn and coauthors. Provable Authority technical paper.

arXiv 2610.22899, version 1, October 2026. Sections 4.2 and 4.3 explain payment integration and enforcement. Existing payment rails need not natively parse ChainIT credentials; integrated controls must govern release and connect execution evidence to authorization.

Technical paper