YourData

Agent-Anchor:

Connecting Differentiated Digital Actors and Systems to Trustworthy Execution

The New Trust Problem Revealed by AI and an Execution Architecture Built on KERI and vLEI

Jaegyeong Yeom, YourData

EXECUTION TRUST ARCHITECTURE 2026

CONTENTS

The New Trust Problem Revealed by AI and an Execution Architecture Built on KERI and vLEI

1. The Problem Is Not Only the Rise of AI, but a Change in How Connections Among Differentiated Actors and Systems Become Action

The rise of AI is an important change. However, the problem addressed by Agent-Anchor does not arise solely from the emergence of a new AI technology.

A more fundamental shift is taking place. Forms of agency, roles and rights, data and assets, control, and the capacity to act—once bundled together within people, organizations, institutions, and systems—are increasingly being expressed as distinct units. These are not all the same kind of thing, but each can now be represented, verified, delegated, controlled, or transferred in ways appropriate to its own nature.

This differentiation does not imply isolation. On the contrary, these newly distinguished and increasingly visible units are reconnecting across the boundaries of organizations, platforms, industries, and jurisdictions. In the process, information and value are exchanged, rights and control are transferred, and new relationships and actions performed on behalf of others are formed.

The development of digital technology has continually changed how such connections are made. Search engines connected distributed information through search and ranking. Social platforms organized people, content, and relationships through social graphs. Marketplaces connected sellers, products, demand, payments, and logistics within a single transaction process. Blockchain and Web3 introduced new ways to jointly record and verify digital assets, contracts, and changes of state.

Today’s shift continues this trajectory, but both what is being connected and the consequences of those connections are becoming more complex. In the near future, an AI agent acting for a company may negotiate commercial terms with a supplier in another country, verify internal policies and budgets, and then proceed through financial, logistics, product-information, and regulatory systems to initiate orders, payments, shipments, and filings in sequence. Human judgment and an organization’s capacity to act may no longer reside only with a single employee or within a single application. They may operate across multiple human representatives, agents, machines, credentials, and enterprise systems.

Recent institutional discussions reflect different dimensions of this transition. NIST characterizes AI systems as inherently socio-technical, while recent analysis published by the World Economic Forum distinguishes an agent’s technical capabilities from the organizational authority it is permitted to exercise.[1,2]

The difficult question in such an environment is not merely whether these systems can be technically connected. It is also necessary to determine whom each agent or system represents, what authority it holds, how another party can verify changes to or the termination of that authority, and whose action and responsibility should attach to an outcome produced across multiple systems.

Nor is this a problem that belongs only to a future that has yet to arrive. Cross-border transactions already involve company employees and ERP systems, carriers and electronic-document systems, financial institutions and trade-finance or payment systems, customs authorities, and regulatory platforms. Across different organizational and jurisdictional boundaries, they perform a sequence of actions involving document issuance, transfers of rights, payments, and regulatory filings.

In this process, invoices, electronic bills of lading, credentials, and supply-chain data are not merely pieces of information being transmitted. They may serve as the basis, object, or result of subsequent actions. The systems involved are likewise no longer only communication channels connecting the parties; they become part of the execution path through which real outcomes are produced.

The defining feature of the current shift, therefore, is not that digital connections have begun to produce economic and institutional consequences for the first time. Rather, an increasing number of actions are being formed and executed through combinations of individuals and legal entities, representatives and agents, systems, authority, data, and digital objects distributed across different organizations, platforms, and jurisdictions.

These are not the same kind of entity. Individuals and legal entities can bear rights, obligations, and responsibility. Representatives and agents may participate in forming or performing actions on their behalf. Systems may mediate or

execute business processes, while data and digital objects may serve as the basis, object, or result of an action. To assess trust in real business activity, these elements should not be collapsed into a single category. Their differences must be preserved while their relationships are considered together.

Through these combinations, money and assets move, documents and regulatory information are submitted, and claims concerning products and supply chains are created or modified. An organization's official intent may be communicated externally, producing approvals, payments, rights, obligations, responsibilities, and regulatory consequences.

AI did not create this structure. It has instead become a catalyst that exposes, more rapidly and explicitly, problems that existing connection models have handled only implicitly. AI governance and agent architecture therefore cannot be designed on the assumption that AI is simply an independent technology or a more intelligent form of software. They must begin with the questions of whom the system represents, which organization, role, and authority it is connected to, and under what conditions it may participate in actions that produce real effects.

These discussions address different parts of the problem; Agent-Anchor treats them as connected dimensions of a broader execution-trust structure.

This requires a form of digital transformation that goes beyond transferring existing procedures into digital systems. It requires redesigning the structure of connection and action itself. In an environment that produces real outcomes, connection cannot stop at discovery, matching, and routing. It must address not only who is participating, but also who is acting on whose behalf, what authority and delegation support the action, what information and policies apply, and to whom its outcome and responsibility can be attributed.

Agent-Anchor is an execution trust architecture for addressing actions performed by AI and other agents within this broader relationship among identity, authority, grounds, execution, and responsibility.

The central question for Agent-Anchor is therefore not simply:

> *"How can AI be authenticated or controlled?"*

The more fundamental questions are:

Who is participating in an action on whose behalf?

To which person, legal entity, organizational role, mandate, or institutional authority can that action be attributed?

How can verifiable identity and authority lead to actual business execution, and how can a counterparty verify the grounds and outcome of that action?

2. AI Has Made This Structure Most Visible

AI did not create these problems.

Long before AI, companies already performed work through employees, external representatives, software, automation systems, suppliers, and platforms. People, organizations, roles, authority, delegation, and data were already distinct from one another.

Existing digital systems, however, were generally designed around identifiable human operators, internal accounts, relatively closed organizational boundaries, and predefined workflows. Questions of who had acted, which organization bore responsibility, and what authority applied could, to some extent, remain implicit within internal IAM systems, approval records, contractual relationships, and organizational procedures.

AI disrupts these assumptions.

AI can interpret natural-language instructions, generate intent and content, produce documents, and interact directly with customers or external systems. It can combine multiple tools and APIs, initiate workflows, and propose actions that may have external effects.

The results and action paths produced by AI may also be less than fully predetermined. Similar circumstances can lead to different judgments, content, or proposed execution paths, and the practical role and impact of the AI may change depending on how it is connected to surrounding data, tools, and systems.

As a result, questions that existing systems had kept largely hidden now come to the surface.

- Whose intent does an AI-generated message represent?
- On behalf of which person or organization is the AI operating?
- Within what role and scope of authority may it use information and tools?
- Does technical access to a system or API mean that it has organizational authority to perform a particular action?
- Can a document, request, payment instruction, or external interaction generated by AI be accepted as an official organizational action?
- To whom should the outcome be attributed, and how can a counterparty verify the grounds for that attribution?

AI cannot therefore be understood merely as one more entity added to an existing digital environment. It extends the judgment and capacity for action of people and organizations into external systems, making it necessary to design the relationships among identity, representation, authority, action, execution, and evidence more explicitly.

The practical role of AI is likewise not determined by model performance alone. Its organizational meaning and risk depend on whom it represents, which current information and credentials it is connected to, which systems and functions it can use, and which outputs it can convert into actual actions.

For this reason, AI trust cannot be treated solely as a matter of model safety or output quality. For AI to participate in real organizational and market activity, its capacity to generate outputs must be distinguished from the conditions under which those outputs can become official organizational actions.

Placing the AI problem within a broader social and organizational structure does not reduce AI to a peripheral example. On the contrary, it provides a more precise account of the conditions and systems required for AI to operate as a trusted participant in real organizations and markets.

3. Interoperability, Policy Autonomy, and Innovation Must Be Designed Together

This problem cannot be solved by integrating every local system into a single centralized global framework. Countries and industries, public authorities and companies, each have their own legal-entity registration systems, regulatory structures, organizational authorities, internal policies, and accountability frameworks. These are not technical inefficiencies to be eliminated, but institutional contexts formed by each society and market.

Accordingly, the role of a global digital organizational identity infrastructure is not to consolidate or replace these local structures. The important task is to preserve the autonomy and context of each system while enabling legal-entity identities, organizational roles, and the grounds for authority and delegation established within them to be recognized and verified across other organizations, industries, platforms, and jurisdictions, and then used within each party's own policy and accountability framework.

From this perspective, interoperability, policy autonomy, and innovation are not competing objectives. They are design conditions that must be accommodated together. Interoperability does not remove local autonomy; it allows trust and authority rooted in local systems to retain their meaning beyond their original boundaries.

GLEIF has framed the relationship among global interoperability, diverse policy priorities, and innovation goals as a potential "Digital Policy Trilemma."[3] Agent-Anchor approaches the same relationship as a design problem rather than an unavoidable trade-off: when trust can travel across boundaries while authority, policy, and execution remain locally governed, interoperability, policy autonomy, and innovation can become mutually reinforcing design conditions. We describe this reframing as a shift from a potential policy trilemma toward a design trinity.

Innovation, likewise, cannot be sustained by bypassing existing institutions and accountability structures. It becomes sustainable when trust can be verified across different institutions and systems and put to practical use within emerging AI-enabled and digital workflows.

4. KERI and vLEI Provide the Trust Foundation for a New Mode of Connection

Under these conditions, it is not sufficient to bring every participant into a single platform account, centralized database, or identity service controlled by a particular vendor. Each entity and organization must be able to preserve its own identifier control and institutional context while allowing other parties to verify the continuity and changes of that control, as well as the relevant delegation relationships.

This requirement follows directly from the structure described earlier. As organizational action becomes distributed across people, AI agents, enterprise systems, credentials, counterparties, and external infrastructures, the environment in which an action is formed can no longer be assumed to reside within a single application or trust boundary. A commercial action may begin with an instruction inside one organization, be interpreted or developed by an agent, checked against internal systems and policies, communicated to another organization, and ultimately produce effects in payment, logistics, document, or regulatory systems.

These systems do not necessarily share a single control state. Each may maintain its own identities, permissions, credentials, data, policies, and records, and each may change independently. An agent may retain technical access while the organizational mandate supporting that access has changed. A credential may remain in circulation after the authority on which it depends has been modified or revoked. A key may rotate while another system still holds an earlier view of control.

The problem is therefore not only how to connect distributed systems. It is also how trust can remain current and verifiable while control, authority, and delegation change across them.

The resulting architecture should not be distributed merely because decentralization is desirable in itself. It must accommodate distribution because the actors, systems, and institutional boundaries through which organizational actions are formed are themselves distributed. Requiring all participants to collapse into a single global account system or control domain would reintroduce a centralized dependency precisely where the underlying environment is becoming more differentiated.

What is needed instead is a way for control to remain local while the evidence required to verify that control can travel across boundaries. A verifier should not need to become part of the controller's internal system in order to determine whether an identifier remains under valid control, whether that control has changed, or whether a relevant delegation still holds.

KERI provides a protocol foundation for this model. It represents identifier control and changes to that control through an authenticated sequence of key events, including key rotation and delegation. Rather than treating an identifier as a static account entry whose current state is defined only by a particular platform, KERI allows the continuity and evolution of control to be verified through its event history.[4]

This has an important architectural consequence. Different organizations and systems can retain their own operational and control boundaries while another party can independently verify the state and history relevant to an identifier and its delegation relationships. Trust does not have to be recreated through a new bilateral account relationship each time an action crosses an organizational or platform boundary.

The significance of KERI therefore lies not only in the individual functions of key rotation, delegation, or event verification. It provides a mode of connection appropriate to an environment in which control remains locally held while actions and relationships extend beyond the place in which that control originates.

Cryptographically verifiable control, however, does not by itself establish the institutional meaning of an identifier. A verifier still needs to determine which real legal entity stands behind an interaction and, where relevant, which organizational role or capacity provides the basis for representation.

GLEIF's vLEI connects the KERI-based verification structure to the institutional trust associated with legal-entity identity and organizational roles. It enables legal-entity identity, together with the organizational capacity of a person who represents that legal entity or performs a specific organizational role, to be presented in a cryptographically verifiable form.[5]

The purpose of this structure is not to replace local legal-entity registration systems, enterprise IAM, sector-specific certification schemes, or national digital identity systems. Those systems remain part of the institutional contexts in which authority and responsibility are established. The role of an interoperable organizational trust layer is instead to allow identities, roles, and relevant trust relationships rooted in those contexts to remain verifiable when interactions and actions move beyond their original boundaries.

In this sense, KERI and vLEI reflect the same structural principle introduced earlier in this paper at a different layer. Differentiated actors and systems do not need to be collapsed into a single system in order to interact. Their differences and local control can be preserved while the relationships required for trustworthy interaction remain independently verifiable.

KERI therefore provides the protocol foundation for verifying control and delegation across distributed boundaries, while vLEI connects that foundation to the institutional trust of real legal entities and organizational roles.

Yet this still does not determine whether a particular action is authorized, whether the applicable policies and approval conditions have been satisfied, or whether the action ultimately executed corresponds to what was approved. These questions require a further operational layer that connects verifiable organizational trust to action-specific authority and execution.

5. From Verifiable Identity and Roles to Action Authority

KERI and vLEI provide an important foundation for verifiable control, legal-entity identity, organizational roles, delegation, and revocation. However, the fact that an agent or system is connected to a particular organization and can use that organization's data and business systems does not automatically make every action it proposes or performs a valid organizational action.

Separate questions remain:

- Does this action fall within the scope of the relevant role or mandate?
- Are the relevant authority and delegation still valid?
- Are the action's target, purpose, and scope—as well as the applicable policies and approval conditions—satisfied?
- Does the action to be executed correspond to what was actually approved?

Technical access and execution capability must therefore be distinguished from organizational authority for a particular action. API permissions, control of keys, or possession of credentials may be necessary conditions, but none of them, by itself, establishes authority for a specific business action.[2]

Verification of legal-entity identity and organizational roles is a necessary starting point, but it does not automatically complete the authorization and execution of a particular action. An operational layer is required to connect verifiable organizational trust to action-specific authority and actual business execution.

6. Agent-Anchor and the Executor

Agent-Anchor is the overall system and execution trust architecture that addresses the gap between verifiable trust and actual action. When a particular agent acts on behalf of a person or organization, or participates in an action on the basis of an organizational role, mandate, or institutional authority, Agent-Anchor addresses the conditions under which that action can become a verifiable and attributable organizational action.

Here, the term agent does not refer only to an AI agent. It includes human representatives, AI agents, automated services, enterprise systems, and robots that form intent or participate in the performance of an action on behalf of a particular person or organization. These are not the same kind of entity, but when they participate in delegated action, they raise a common class of trust questions.

Agent-Anchor distinguishes between the Agent and the Actor. An Agent is an operational participant that proposes or performs an action. An Actor is the person or legal entity that serves as the reference point for determining the organizational and legal attribution of that action.

The fact that an Agent can generate or technically perform an action does not, by itself, make that action an official organizational action. The action must be connected to verifiable legal-entity identity, organizational roles, mandates, authority, and policy.

Within Agent-Anchor, the Executor is the core runtime that handles this connection. If Agent-Anchor refers to the overall architecture, the Executor is the execution layer that connects verifiable organizational trust to the authorization, execution, and evidence associated with a particular action.

This distinction should not be understood as requiring the Agent and the Executor to exist as isolated systems or in completely separate operational environments. Agent-Anchor may be integrated with the broader environment in which agents receive context, use tools, interact with enterprise systems, and participate in workflows. What must remain distinct is the basis on which an action is formed from the basis on which that action is permitted to produce an organizational effect. The fact that an Agent generated, selected, or technically initiated an action cannot itself constitute the authority for that action.

The Executor therefore represents a logically distinct authority and execution boundary within this broader environment. Its role is not to separate trust from the agent ecosystem, but to make the relationship among organizational identity, authority, policy, the proposed action, and consequential execution explicit and enforceable. This allows Agent-Anchor to participate in an integrated agent and enterprise environment without collapsing the generation of an action and the authorization of that action into the same decision.

The Executor is neither a mere credential verifier nor a reasoning component inside the AI. It evaluates the organizational attribution of an action, the basis of authority supporting it, the applicable policies, and the relevant execution conditions. It then ensures that only permitted actions are reflected in actual business systems and external environments.

KERI and vLEI provide the trust foundation required by this structure: verifiable control, legal-entity identity, and organizational roles. Agent-Anchor connects this trust to the authorization and execution of concrete actions.

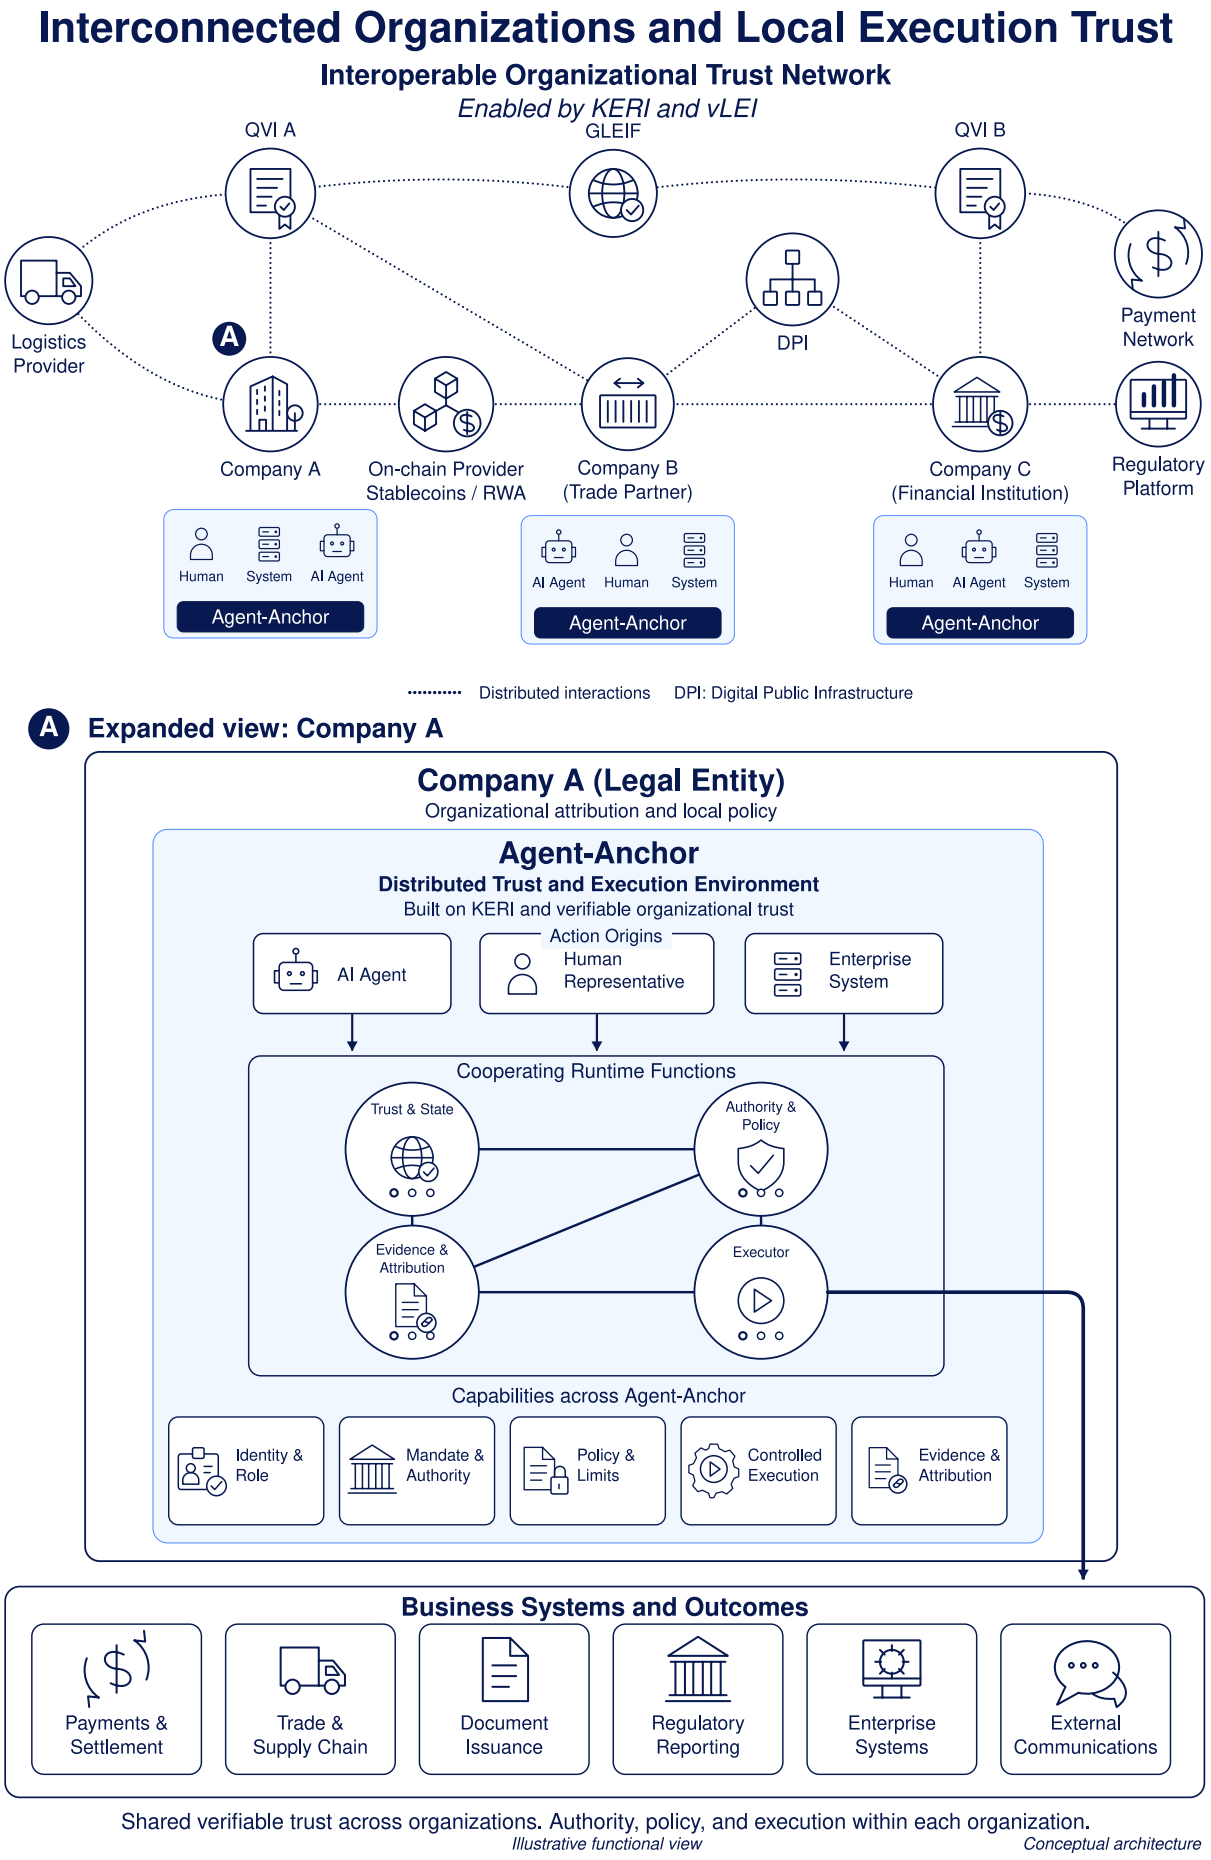

Figure 1. Interconnected Organizations and Local Execution Trust

Agent-Anchor connects interoperable organizational trust across entities to locally governed authority, policy, execution, and evidence within each organization.

7. How Agent-Anchor Addresses AI

AI is not the entirety of the problem addressed by Agent-Anchor, but it is one of the clearest and most urgent cases in which this structure is needed. AI can generate intent and content that are not fully predetermined and, depending on the circumstances, combine multiple tools and systems to propose actions with external effects.

Agent-Anchor does not seek to eliminate this non-deterministic generative capacity or to control every part of the AI's internal reasoning. Instead, it distinguishes between the domain in which AI generates outputs and actions, and the authority boundary through which those outputs may become official organizational actions.

As agents become connected to more information, tools, memory, and external systems, the path through which an action is formed may itself span multiple components and change dynamically according to the agent's observations and decisions. Agent-Anchor does not require this entire process to become deterministic. Instead, it requires the transition from a proposed action to an action with organizational consequences to remain subject to explicit, verifiable, and enforceable conditions.

This allows the agent environment and the execution-trust environment to be integrated without treating them as identical. An agent may retain flexibility in how it interprets an objective, gathers information, selects tools, or forms a proposed action, while the authority under which that action may affect an external system remains independently assessable.

The fact that AI can generate—or technically perform—a message, document, request, or other action does not automatically make the result an official action of an organization. Before an action with external effects is executed, it must be made clear, in a verifiable form, what is to be performed, on whose behalf, and under what authority and conditions.

The Executor connects the proposed action to legal-entity identity, organizational roles, mandates, applicable policies, and execution conditions. Only actions that satisfy the required conditions are reflected in actual business systems or external environments, and the grounds of authority and the resulting execution may be connected to verifiable evidence.

This structure preserves AI's flexible generative capacity while requiring actions with external effects to pass through a verifiable and enforceable organizational authority boundary.

8. The Boundary Between Internal AI Controls and Execution Trust

The fact that Agent-Anchor addresses an important dimension of AI trust does not mean that it directly solves every AI safety problem. Model training methods and weights, alignment, the quality of internal reasoning, hallucination mitigation, and safety training at the foundation-model level remain the responsibility of model providers and dedicated AI safety and security systems.

At the same time, Agent-Anchor is neither a simple identity layer unrelated to those systems nor a gateway that examines an output only at the final stage after generation is complete. Agent-Anchor addresses whom an agent represents, the context of organizational trust and authority in which it uses information, tools, and systems, and the business paths and actions in which it may participate.

Existing controls—including guardrails, AI firewalls, data loss prevention systems, and API or tool gateways—can detect risks or enforce conditions of use within their respective domains. Rather than replacing them, Agent-Anchor can connect these controls to the relevant organizational identity, role, mandate, purpose, and action context, and incorporate their determinations into the authorization of an actual action.

The objective is not to place every control function inside a single product. It is to enable identity, authority, the conditions governing the use of data and tools, execution, and evidence—which have often been managed separately—to operate within a single verifiable action context.

Agent-Anchor therefore does not directly control the entirety of an AI system’s internal reasoning process. Instead, it establishes an execution environment in which an agent operates within an organization’s trust and authority structure, while applying a verifiable execution boundary to actions that produce external effects.

9. The Executor Connects vLEI Trust to Actual Execution

GLEIF and vLEI make legal-entity identity and organizational roles verifiable.[5] However, neither GLEIF nor an individual QVI can directly define and execute every specific business action, policy, and workflow across all industries and organizations.

The organizational trust provided by vLEI can be used across many different environments, but the authority, policies, risk conditions, and methods of execution applicable to a particular action vary by domain. For verifiable organizational trust to be put into practical use, an operational layer is therefore needed to connect it to action-specific authority and execution within each business environment.

The Executor performs this role. It connects verifiable trust in legal-entity identity and organizational roles to the mandate, policies, and execution conditions applicable to a particular action. It also enables permitted actions to be reflected in actual business systems, while allowing the grounds for authorization and the resulting execution to be connected to verifiable evidence.

This does not mean that the Executor replaces the policies or workflows of every domain with a single standardized model. It preserves the rules and execution environments specific to each industry and organization while providing a common structure through which verifiable organizational trust can be translated into action.

Agent-Anchor is therefore not merely a vertical solution that applies vLEI to one particular use case. It proposes a reusable execution model that connects the trust expressed and verified through vLEI to actions governed by the authority, policy, and execution conditions of different domains.

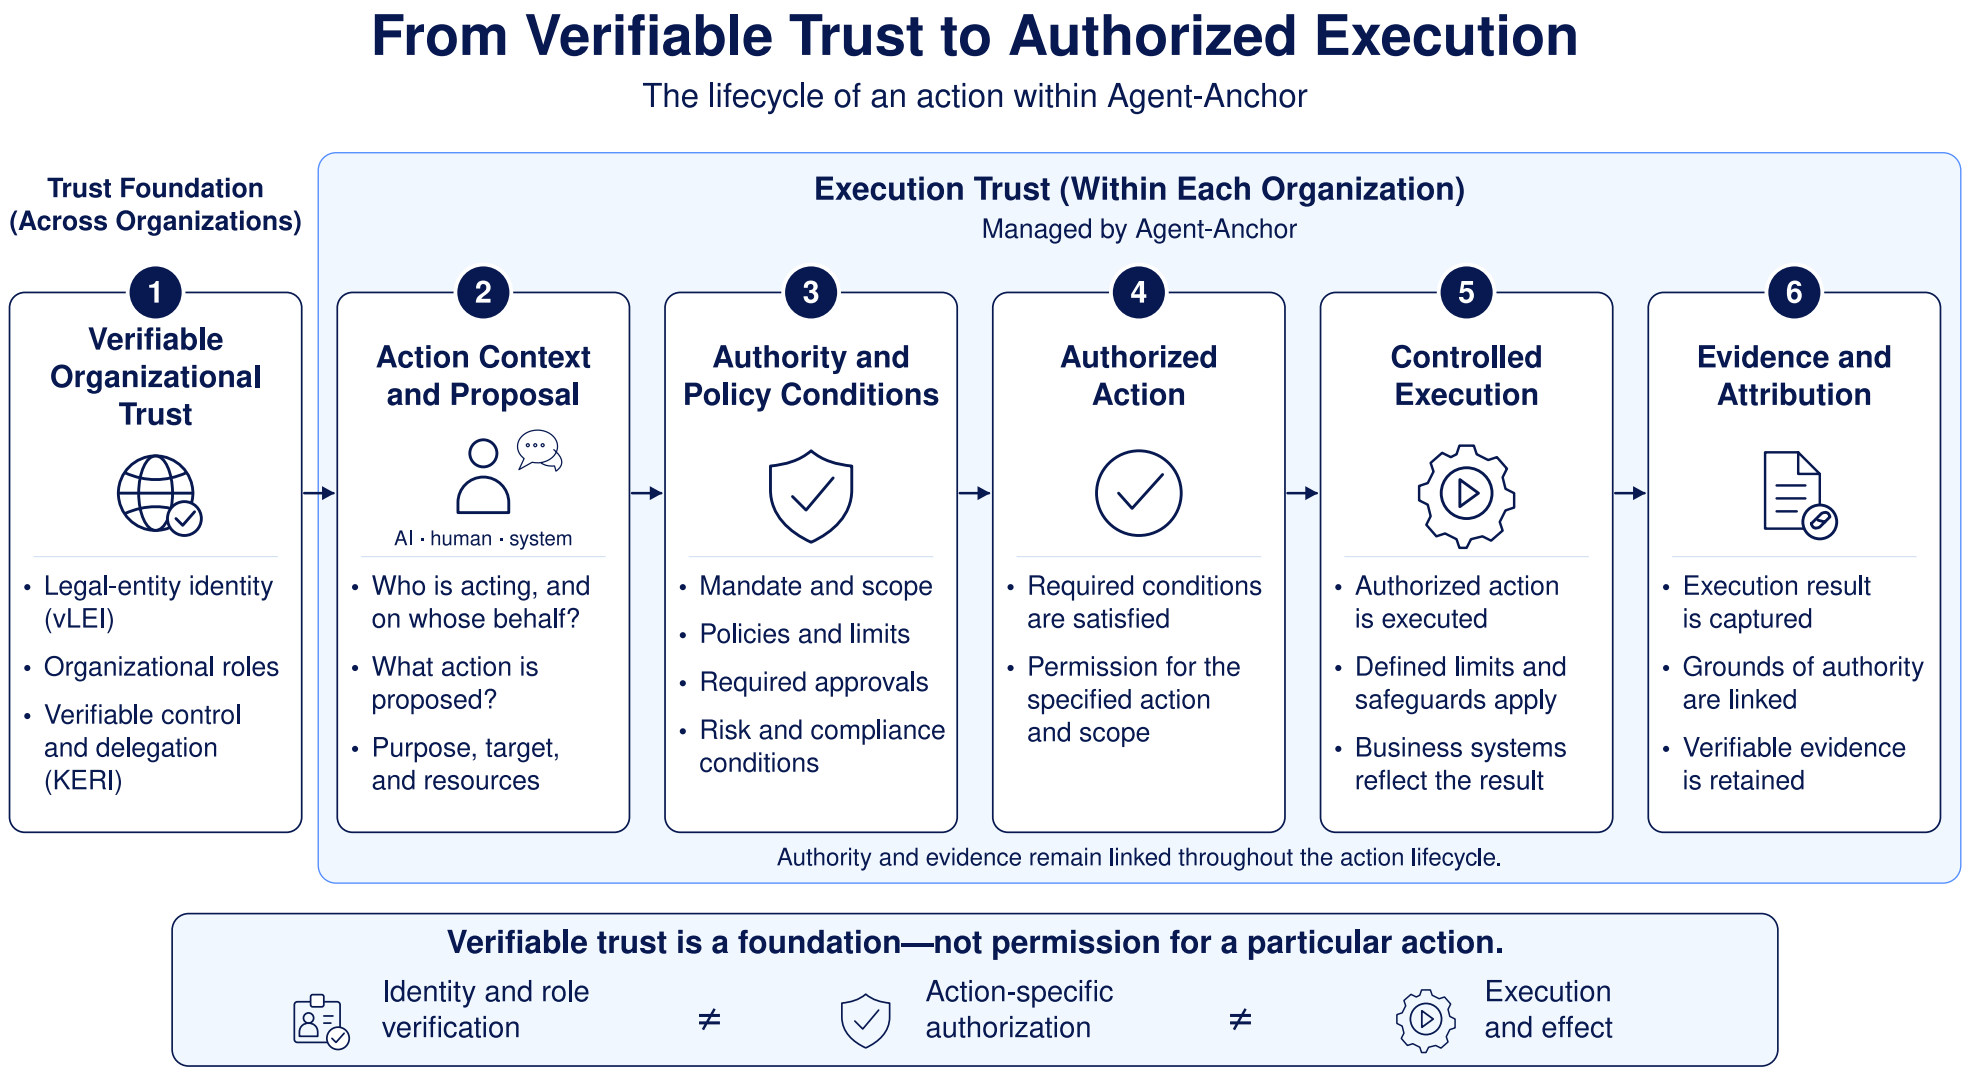

Conceptual action flow; not a component topology or detailed protocol sequence.

Figure 2. From Verifiable Trust to Authorized Execution

Agent-Anchor connects verifiable organizational trust and action context to action-specific authorization, controlled execution, and attributable evidence within each organization.

10. The Organizational Trust Anchor Provided by vLEI

In an environment where people, AI agents, and automated systems participate in actions on behalf of organizations across organizational, platform, and jurisdictional boundaries, it must be possible to determine which legal entity and organizational role provide the basis for a particular action.

vLEI provides an institutional reference point for this determination by making legal-entity identity and organizational roles verifiable.[5,6] Its significance does not lie in merely attaching an organization’s name to an agent or system. It enables another party to verify the legal entity and organizational capacity to which a person, agent, or system is connected.

This provides a verifiable basis for assessing the organizational representation and attribution of delegated actions. However, verification of legal-entity identity and organizational roles does not, by itself, automatically determine the authority or validity of a specific action, or the final allocation of responsibility.

Within Agent-Anchor, vLEI serves as a core trust anchor for connecting an action to the legal entity and organizational structure that stand behind it. The Executor connects this organizational trust to the relevant mandate, policies, and action context so that it can be applied in actual business operations.

vLEI therefore does not define the entire scope of Agent-Anchor as a single feature. It is, however, a core foundation of organizational trust that enables actions performed on behalf of organizations by people, AI, and other systems to be connected to verifiable legal-entity identities and organizational roles.

11. The Same Structure Repeats Across Different Use Cases

The initial use case for Agent-Anchor involved AI callbots, chatbots, and automated messaging systems that interact with customers on behalf of companies.

In this context, customers need to verify more than the authenticity of a voice or phone number. They must be able to determine whether the AI actually represents the company, under which organizational role and authority it operates, and whether an action proposed by that AI can be accepted as an official organizational action.

The same problem recurs across other domains.

In a Digital Product Passport environment, it must be possible to verify who created or modified a product or compliance claim, on the basis of which legal entity, organizational role, and mandate, and whether the change fell within the permitted scope.

EU law is already formalizing the Digital Product Passport as a regulated information structure: under the Ecodesign for Sustainable Products Regulation, passport data must be accurate, complete, and up to date.[7] Agent-Anchor asks a different but complementary question: which actor was authorized to create or modify a consequential claim within that structure?

In payments, an AI or enterprise system may be capable of generating a payment request, but the authority and approval conditions applicable to the amount, recipient, and purpose of the transaction must still be satisfied separately.

Related financial-system analysis exposes the same separation from another direction. An IMF Note on agentic payments distinguishes intent, authorization, and settlement, while a 2026 Bank of Japan Review describes AI-agent commerce through sequential mandates that confirm user intent, transaction conditions, and payment execution.[6,8] Agent-Anchor generalizes this distinction beyond payments: generating or initiating an action is not the same as holding the authority under which that action may become an official organizational act.

In trade and supply chains, verifying the organization from which a document originated is not sufficient. It is also necessary to determine which role issued the document, under what authority, and whether its actual content corresponds to the action that was approved.

That distinction becomes concrete in electronic transferable records. UNCITRAL’s Model Law on Electronic Transferable Records treats control as the functional equivalent of possession for electronic transferable records such as bills of lading, making control legally consequential rather than merely informational.[9]

The same distinction applies to corporate wallets and digital assets. The ability to control a key and sign a transaction does not, by itself, justify every asset-related action. Key control must be connected to legal-entity identity, organizational roles, mandates, transaction scope, and applicable policies.

These developments do not define a single architecture and arise from different domains. Taken together, however, they expose a recurring pattern: consequential digital activity increasingly depends on connecting information and technical capability to authority, control, and verifiable execution. Agent-Anchor treats that recurring pattern as a common execution-trust structure.

Across these different industries and workflows, the same underlying structure recurs:

An Agent or system proposes an action.
→ Verifiable organizational trust and the grounds of authority are presented.
→ The scope of the action and the applicable conditions are assessed.
→ Only permitted actions are reflected in the actual business environment.
→ The underlying grounds and execution results are connected to verifiable evidence.

12. YourData’s Perspective: Extending AI Trust into Execution Trust

YourData does not view Agent-Anchor as a single-function product that merely assigns identities to AI agents or blocks particular risks. Agent-Anchor is an execution trust architecture for an environment in which the judgment, information, authority, and capacity for action of people and organizations are extended through human representatives, AI, machines, and software systems, and in which these actors and systems produce real outcomes across boundaries.

In the future, a single organizational action may no longer be completed by one employee or within one application. A human-defined objective, organizational policy, intent and plans formed by AI, approvals associated with different roles, and execution through external systems may together form a single action chain across different places and points in time.

In this environment, AI may become a cognitive and operational extension through which people and organizations perceive and interpret information, interact with other systems, and perform work. The role that AI can perform in practice is not determined by model performance alone. It also depends on whom the AI represents, which current information, credentials, authority, and systems it is connected to, and which results it can convert into actual action.

The central problem of AI trust is therefore not limited to how intelligent AI can become. Even when the judgment and capacity for action of people and organizations are distributed across multiple agents and systems, it must remain possible to explain and verify who is acting on whose behalf, what information and authority support the action, when that action becomes an official organizational action, and to whom its outcome can be attributed.

Whether Agent-Anchor applies does not depend on whether a particular system is formally described as an “AI agent.” The same execution trust problem arises when existing automation or a process combining people and machines uses information, proposes actions, or produces external effects on behalf of another party.

Agent-Anchor does not design the training process or internal generative mechanisms of a foundation model. Nor is it limited to serving as a gateway that examines an output only at the final stage. It forms an execution environment in which an Agent can use information and capabilities and participate in business processes within an organization’s trust and authority structure. It also applies a verifiable authority boundary to actions that produce external effects.

KERI provides a foundation through which entities and organizations can retain control while enabling the continuity of that control, changes to it, and delegation to be verified across boundaries. vLEI connects this foundation to the institutional trust of legal-entity identity and organizational roles. The Agent-Anchor Executor connects this organizational trust to the mandates, policies, and execution conditions applicable to a particular action, allowing it to operate within actual business processes and to be reflected in verifiable evidence.

Agent-Anchor is therefore neither a simple AI identity layer nor merely an individual guardrail or execution gateway. It is an execution trust architecture that enables societies and organizations to adopt AI as a cognitive and operational extension without losing sight of whom the AI represents, which information and authority it is connected to, and what actions it performs.

Agent-Anchor does not seek to manage humans and AI as two completely separate categories of existence. Its objective is to preserve the connection among identity, authority, action, outcomes, and responsibility even when the judgment and capacity for action of people and organizations are distributed across AI, machines, agents, and multiple systems, and when these participants act on behalf of one another across boundaries.

This is a challenge at a different level from making AI more intelligent. It is also one of the fundamental trust conditions required for societies and organizations to adopt AI and machines as real components of their operations and economic activity.

References

[1] Tabassi, E. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1, National Institute of Standards and Technology, 2023. https://doi.org/10.6028/NIST.AI.100-1

[2] Duet, A. “Why Autonomous AI Demands a New Model of Organizational Authority.” World Economic Forum, May 26, 2026. https://www.weforum.org/stories/all/autonomous-ai-new-model-organizational-authority/

[3] Global Legal Entity Identifier Foundation (GLEIF). “Call for Papers: The Digital Policy Trilemma.” 2026. https://www.gleif.org/en/organizational-identity/research-publications/call-for-papers-the-digital-policy-trilemma

[4] Trust over IP Foundation. KERI Specification. https://trustoverip.github.io/kswg-keri-specification/

[5] Global Legal Entity Identifier Foundation (GLEIF). vLEI Ecosystem Governance Framework, Version 4.0, 2026. https://www.gleif.org/en/organizational-identity/become-a-vlei-issuer-qvi/vlei-ecosystem-governance-framework

[6] Yamada, N., Shinzaki, T., Shimizu, T., and Okabe, K. “Verifiable Credentials for Identity Assurance in the Digital Society: An Overview and Trends in Standards Development.” Bank of Japan Review 2026-E-6, April 2026. https://www.boj.or.jp/en/research/wps_rev/rev_2026/data/rev26e06.pdf

[7] Regulation (EU) 2024/1781 establishing a framework for the setting of ecodesign requirements for sustainable products, Chapter III (Digital Product Passport), 2024. https://eur-lex.europa.eu/eli/reg/2024/1781/oj

[8] Davidovic, S., and Tourpe, H. “How Agentic AI Will Reshape Payments.” IMF Notes 2026/004, April 2026. https://doi.org/10.5089/9781513533308.068

[9] United Nations Commission on International Trade Law (UNCITRAL). Model Law on Electronic Transferable Records, 2017. https://uncitral.un.org/en/texts/ecommerce/modellaw/electronic_transferable_records