Digital Identity for Agentic Systems

Two Enterprise Use Case Explorations

Decentralized Identity Foundation Special Report, December 2025

Research Lead: Damian Glover, Independent

Executive Brief

Enterprises are moving from AI copilots toward more autonomous, agent-based systems capable of executing workflows, negotiating terms, and making decisions with limited human oversight. While most deployments today remain within organizational boundaries, pressure is growing to extend these systems across suppliers, partners, regulators, and customers.

As agents begin to act across organizational and legal boundaries, new trust challenges emerge. Enterprises must establish not only *who* an agent represents, but *what authority it has*, *under what conditions it may act*, and *how its actions can be audited and revoked*. Existing identity and access models, designed for user login and synchronous API access, struggle to support these requirements.

This brief explores two representative enterprise use cases**—insurance claims processing** and **regulatory reporting and compliance—**to surface the identity and trust requirements that emerge when autonomous agents operate across organizational boundaries. These use cases are not intended to be exhaustive or prescriptive; rather, they provide concrete lenses through which to understand a broader architectural shift.

Across both use cases, a common pattern appears. Agent-based workflows increasingly involve asynchronous interactions among multiple parties, each operating under distinct legal, contractual, and regulatory constraints. Authority must be delegated to non-human actors in a way that is explicit, limited in scope, and revocable. Decisions made autonomously must be traceable, auditable, and attributable to responsible organizations.

Current approaches rely heavily on centralized platforms, API integrations, and bearer-token authorization models. While effective for tightly coupled systems, these approaches introduce friction and risk as autonomy increases and interactions span organizational boundaries. They offer limited support for durable delegation, end-to-end auditability, and accountability when autonomous actions have legal or financial consequences.

Across both use cases, five identity-related requirements consistently emerge:

1. Delegated authority for non-human actors that is explicit, scoped, and revocable.
2. Verifiable agent identity across organizational boundaries.
3. Policy-bound authorization rather than broad, bearer-based access.
4. End-to-end auditability of autonomous decisions and actions.
5. Clear accountability mechanisms when agents act on behalf of organizations or individuals.

Taken together, these requirements point to a shift in the role of digital identity. Identity is no longer solely about authenticating users or applications at the moment of access. It must also support ongoing, delegated authority across autonomous interactions that unfold over time and across trust boundaries.

Why Cross-Boundary Agent Use Cases Matter

Early enterprise AI deployments focus on copilots and internal automation, where trust boundaries are largely implicit and controlled. As organizations move toward autonomous agents capable of acting independently, these systems increasingly interact with external parties: customers, partners, suppliers, and regulators.

Once agents cross organizational boundaries, trust assumptions change. Actions may carry legal, financial, or regulatory consequences, and failures can propagate across ecosystems rather than remaining contained within a single system. Identity, authority, and accountability become explicit design requirements rather than implementation details.

These dynamics make cross-boundary use cases an effective lens for understanding where existing identity and access models begin to break down in an agentic world.

How to Read the Use Cases

The use cases that follow are exploratory rather than prescriptive. Each uses a common structure to make patterns easy to compare across domains.

The focus is not on specific technologies or vendors, but on:

- where trust boundaries appear,
- what agents actually do,
- and which identity capabilities are required for safe operation at scale.

Use Case 1: Agentic Insurance Claims Processing

Business Context & Trigger

Insurance claims processing is manual, fragmented, often slow and susceptible to fraud. The associated administrative expenses, loss adjustment costs and fraud losses form a significant portion of insurance costs.

Advances in AI, particularly in document understanding, decision support, and workflow automation, have demonstrated that large portions of claims handling can be executed autonomously, at least for straightforward cases. Meanwhile, customer expectations are shifting, driven by digitally native insurers that market rapid, automated payouts.

While current agentic initiatives tend to remain internal, the economic logic points toward distributed, agent-to-agent interaction across the insurance ecosystem.

Actors & Trust Boundaries

The main actors are policyholders, insurers, loss adjustors, third-party service providers (such as repair shops or medical providers), and, for larger exposures, reinsurers. Regulators sit outside the transaction flow but impose constraints on data handling, decision transparency and governance.

In an agentic model, each of these parties may be represented by one or more software agents acting on their behalf. A customer-facing agent embedded in a mobile application may collect evidence and submit a claim. An insurer-operated agent may validate coverage, route the case, and coordinate downstream activities. External providers may expose agents that submit estimates or bills. Reinsurers may operate agents that verify eligibility and exposure thresholds.

Trust boundaries arise wherever authority, information, or liability crosses from one organization to another. These boundaries are not purely technical (such as data entering or exiting an entity's IT environment). They also reflect contractual relationships, regulatory jurisdictions, and accountability regimes.

Agent Interactions (What Actually Happens)

A representative workflow begins when an automotive policyholder (or their agent) presents a claim. The insurer's claims agent verifies the policyholder's identity, confirms the policy is in force, assesses the incident details and evidence, and determines the claim category. The agent

may be authorized to approve resolution directly for low-risk claims. More complex cases are routed to specialist agents or escalated to a human.

Downstream interactions quickly extend beyond the insurer's perimeter. An adjuster agent requests repair estimates from a network of service providers, each potentially represented by its own agent capable of proving licensing, insurance coverage, and prior performance. In injury-related claims, medical provider agents submit treatment and billing information with selective disclosure, sharing only details relevant to the claim.

A reinsurance agent is automatically notified if the claim exceeds predefined thresholds. It verifies that the primary insurer's agent is acting within its delegated authority and confirms coverage before acknowledging liability. Finally, once all conditions are met, a payment agent executes disbursement and records an immutable audit trail capturing the sequence of decisions and authorizations.

Humans are still involved, but their role shifts from execution to oversight. Rather than manually processing every step, they supervise exceptions, review edge cases, and define policies that govern agent behaviour.

Trust & Identity Requirements

Identity and trust capabilities are foundational for this workflow to function safely at scale. Unlike traditional API integrations, where access is tightly pre-negotiated and centrally managed, agent interactions are contextual and decision-bearing. Each party must be able to determine not only that an agent is authenticated, but also whom it represents, what authority it holds, and under what constraints it is operating. For such a system to be trustworthy, all of the following conditions must be met.

1. Each agent needs a verifiable identity that persists across organizational boundaries.
2. Authority must be explicitly delegated. An agent approving a payout or requesting sensitive data must carry cryptographically verifiable proof of what it is allowed to do, within what limits, and for how long.
3. Policy enforcement must be intrinsic to agent operation. Rules governing data minimization, confidentiality, jurisdictional compliance, and escalation thresholds need to travel with the agent's authority, not be enforced solely by a central platform. This is particularly critical when handling regulated data such as medical information.
4. The entire process must be auditable end to end. Autonomous decisions must be traceable to specific agents, delegations, and input evidence, not only for regulatory compliance, but also for dispute resolution and liability management.
5. Delegation must be revocable. If an agent is compromised, misconfigured, or exceeds its mandate, its authority must be withdrawn without dismantling the entire system.
6. Accountability must be clear. When errors occur, responsibility must be attributable to the organization that empowered the agent.

What Breaks with Today's Approaches

Current claims platforms rely on centralized workflow engines, point-to-point integrations, and bearer-token authorization models. These approaches work when interactions are synchronous, short-lived, and confined within a single enterprise or tightly governed partner network.

As autonomy increases, these models break down. Bearer tokens provide access but carry little contextual information about delegated authority or intent. Platform-mediated integrations create brittle dependencies and limit the ability to reason about actions that unfold over time. Manual controls and after-the-fact audits are unable to keep pace with real-time, agent-driven decisions.

Most critically, existing models conflate authentication with authorization. They can confirm that a system is allowed to connect, but not whether a specific autonomous action is appropriate. This gap introduces risk where insurers are most sensitive: fraud exposure, regulatory compliance, and customer trust.

Feasibility & Adoption Horizon

Autonomous claims agents are on the industry's radar. Major European insurers are expected to commence research and development programs in 2027, with end-to-end claims processing automation anticipated to be in production by 2032.

Elements of the use case described above are feasible today. Insurers already deploy automated decisioning for simple claims and are experimenting with AI-assisted adjuster workflows. Extending these capabilities to structured, verifiable agent interactions with external providers is a logical next step, especially where strong commercial relationships already exist.

Insurer-issued Verifiable Credentials (VCs) are already beginning to be used in cross-boundary contexts: partnerships between insurers and logistics ecosystems have demonstrated how proof-of-insurance VCs can be instantly validated by third parties without the need for bilateral integrations, to automate acceptance or compliance checks.

The insurance sector also has experience with cross-organizational automation through an earlier consortium focused on blockchain-based reinsurance and contract reconciliation.

These efforts show that the industry is not starting from zero; rather, it is incrementally assembling the identity, trust, and governance components required for agentic interaction.

Legacy systems, fragmented data, and the absence of common messaging and semantic standards remain significant barriers. Alignment with interoperable trust infrastructure such as the EUDI Wallet—which EU insurers are mandated to accept for customer authentication by

December 2027—European Business Wallet (EUBW) and Verifiable Legal Entity Identifier (vLEI) can help bypass these hurdles and bridge towards cross-organizational agent ecosystems.

Regulatory engagement is essential, as supervisors are increasingly scrutinizing insurer AI governance and may require standardized reporting and controls.

Use Case 2: Agentic Regulatory Compliance Across Financial and Trade Networks

Business Context and Trigger

Regulatory compliance has traditionally been implemented through batch reporting, portals, and manual attestations. As economic activity becomes faster, more distributed, and increasingly automated, this model is under strain.

Compliance is becoming simultaneously more demanding and less compatible with existing approaches. Financial institutions face escalating fraud volumes driven by AI-enabled criminal activity, while regulators impose stricter obligations for real-time detection and consumer protection. In parallel, global trade flows involve trillions of dollars in daily value, yet customs and import processes remain document-heavy and slow.

In both domains, the core challenge is timing and scale. Fraud must be detected before a payment completes, not weeks later through reconciliation. Import filings must be accurate and verifiable at the moment goods cross borders, not retrospectively assembled during audits.

These pressures are pushing regulators and industry toward continuous, machine-mediated compliance rather than episodic reporting. Two emerging trends illustrate this shift: cross-institution fraud prevention in financial services, and automated trade settlement and import compliance in global supply chains.

This shift is emerging now for three reasons. First, AI systems are increasingly capable of extracting, classifying, and reasoning over complex regulatory data. Second, privacy and data-protection constraints make centralized data pooling untenable, forcing exploration of distributed and privacy-preserving approaches. Third, regulators themselves are experimenting with automated verification and machine-consumable filings, creating an opening for agent-based interaction rather than human-facing portals.

Actors and Trust Boundaries

The actors in financial fraud prevention include banks, payment service providers, clearing networks, and regulators. In trade compliance, actors include exporters, importers, logistics providers, financial institutions, and customs authorities.

Each organization may deploy one or more software agents to fulfill compliance-related responsibilities. A bank may operate a fraud-detection agent that evaluates transactions in real

time. A payment network may coordinate shared intelligence across participants. A customs authority may operate agents that validate import filings against regulatory rules and risk profiles.

Trust boundaries are pervasive. Institutions are prohibited from freely sharing raw customer or transaction data, yet are expected to collaborate to reduce systemic risk. Regulators require assurance that automated decisions comply with law, but cannot feasibly inspect every transaction manually. As a result, agents must be able to interact across organizational and jurisdictional boundaries while carrying verifiable proof of identity, authority, and policy constraints.

Agent Interactions (What Actually Happens)

In a shared fraud-signals scenario, each participating financial institution operates a local fraud-detection agent embedded within its transaction processing systems. As a payment is initiated, the agent evaluates behavioral and contextual signals in real time. Rather than sharing raw transaction data externally, the agent generates privacy-preserving fraud indicators.

These indicators are exchanged with peer institutions through a federated network. Each institution is a node, training models locally on its own data. Encrypted or aggregated updates are shared to strengthen collective detection capabilities without exposing sensitive information. A coordinating agent aggregates these signals to identify suspicious accounts or transaction patterns that would not be visible to any single institution.

A concrete example of this architecture is already being trialed. In September 2025, a global payments network coordinated a pilot involving multiple major banks, combining AI, federated learning, and privacy-enhancing technologies to share fraud-related signals across borders. In testing on millions of synthetic transactions, the collaborative model reportedly doubled fraud-detection effectiveness compared to institution-specific models, while avoiding direct data sharing. This illustrates how institutions can function as autonomous yet cooperative agents in a shared compliance network.

In automated trade compliance, AI agents representing exporters or importers extract shipment, contract, and provenance data directly from operational systems, assemble machine-readable import filings and submit them to customs authorities without human intervention.

Agents operated by customs authorities automatically verify submissions against tariff schedules, licensing requirements, and risk indicators. Linked data models allow authorities to trace goods through complex supply chains with granular precision. In pilots associated with government innovation programs, this approach has delivered dramatic efficiency gains, including substantial reductions in administrative costs for importers and near-instant document retrieval during investigations.

Across both domains, humans remain responsible for defining policies, thresholds, and escalation rules. Agents execute compliance continuously, while humans intervene only when anomalies or disputes arise.

Trust & Identity Requirements

First, each agent must have a verifiable, non-human identity that counterparties and regulators can authenticate across organizational boundaries.

Second, authority must be explicitly delegated and machine-verifiable. A fraud-detection agent may be authorized to share specific categories of anonymized signals, but not raw customer data. A trade compliance agent may be authorized to submit filings for a defined set of goods, jurisdictions, or time periods.

Third, policy enforcement must be embedded within agent operation. Privacy constraints, data-retention rules, and jurisdictional limitations must govern what agents can share or disclose. The use of privacy-enhancing technologies and federated learning in existing fraud pilots demonstrates how policy can be enforced technically, rather than relying solely on contractual assurances.

Fourth, auditability and traceability are mandatory. Regulators must be able to reconstruct how a decision or submission was produced, which agents contributed signals, and under what authority.

Finally, revocation and accountability mechanisms must exist. If an agent behaves incorrectly or a model is found to be biased or flawed, its authority must be withdrawn without disrupting the broader network. Responsibility must remain clearly attributable to the organization that deployed and governed the agent.

What Breaks with Today’s Approaches

Traditional compliance models rely on centralized reporting systems, periodic disclosures, and manual review. These approaches assume that compliance can lag operational reality and that data can be consolidated into regulatory silos. In real-time fraud prevention and high-velocity trade, these assumptions no longer hold.

Bearer-token authorization and static integrations provide access but fail to express fine-grained, contextual authority. Portal-based submissions introduce latency and cost, and they scale poorly as transaction volumes grow. Existing models also struggle to support collaborative compliance, where multiple institutions must contribute signals or attestations without exposing proprietary or personal data.

The result is a growing gap between regulatory expectations and technical capability.

Feasibility & Adoption Horizon

Elements of agentic regulatory compliance are already implementable and, in some cases, operational. Privacy-preserving fraud-signal sharing and federated learning have moved beyond theory into live trials. Automated import filings and machine verification are delivering measurable cost and time savings in government-backed pilots.

Near-term adoption is most likely in bounded networks where participants share strong incentives and governance frameworks, such as payment networks or specific trade corridors. Wider adoption will require common protocols, semantic standards, and clearer regulatory guidance on non-human delegated authority.

In the longer term, as regulators become comfortable consuming machine-generated attestations and as identity infrastructure matures, these agent-based models could become the default mechanism for high-volume, cross-border compliance.

Cross-Cutting Identity Implications

Similar identity requirements recur across both agentic insurance claims processing and agentic regulatory compliance use cases:

- **Delegated authority:** Agents require explicit, limited authority tied to purpose and context.
- **Verifiable representation:** Agents must be able to prove who or what they represent across organizations.
- **Policy-bound access:** Authorization must reflect business and regulatory rules, not just technical access.
- **Auditability:** Autonomous actions must be traceable and reviewable after the fact.
- **Revocation and accountability:** Authority must be revocable, and outcomes attributable to responsible parties.

These requirements exceed the capabilities of many existing identity and access management systems, particularly as autonomy and cross-boundary interactions increase.

Conclusion

Traditional identity and access models were designed for human users and stateless applications, not autonomous actors operating asynchronously across trust boundaries.

The two use cases presented here highlight a fundamental shift in the role of digital identity. Identity is no longer just about granting access to systems; it becomes the mechanism by which authority is delegated, constrained, and audited over time.

In agentic insurance claims processing, identity must bind together the agent, the organization it represents, the policies it enforces, and the decisions it makes. In agentic fraud detection and trade compliance, identity must provide granular traceability and accountability to enable automatic compliance verification and risk reduction in real time.

Pilot projects show the demand for, and feasibility of, more advanced identity models. What is missing today are mature capabilities for verifiable agent identity, durable delegation and policy-bound authorization that can travel across organizational lines.

As enterprises deploy more autonomous systems, digital identity becomes foundational infrastructure rather than a supporting service. The use cases in this brief illustrate why existing models are under strain and why evolving approaches to identity, delegation, and trust will be essential to scaling agentic systems safely and responsibly.