SignedBy

WORKING PAPER · AUGUST 2026

Agent-to-Agent Signing in Europe

What eIDAS, the Legal Context Protocol, and current practice actually establish — and the real gap that's still open.

SignedBy — signedby.ai

A trading name of SPRK10 B.V.

ABSTRACT

This paper is not legal advice, and it is not an announcement. It is an honest look at where the law and the emerging standards landscape actually stand on AI agents making binding commitments on a company's behalf, written from the perspective of a European e-signature provider with a genuine, live piece of relevant infrastructure and no pretense of having solved the hard part. Two real anchors exist today: the EU's decade-old seal/signature distinction, and a new industry standard for discovering legal terms. Neither one, alone or together, resolves how two independent organizations' agents actually form a binding agreement with each other. That gap is the honest subject of this paper.

Over the past year, AI agents have gone from answering questions to taking actions — booking, purchasing, negotiating, and increasingly, agreeing to terms. The technology has moved faster than the legal framing around it. A narrower, more concrete question sits underneath the broader "can an AI agent sign a contract" debate: when an autonomous agent commits an organization to something, who is legally speaking, under what authority, and how would a counterparty verify any of that? In the United States, that question has decades of case law behind it. In Europe, it has a purpose-built legal instrument that predates the current moment by ten years — and a live gap that instrument was never designed to close.

1. Two real anchors — not one

The instrument Europe already has: seals, not signatures

eIDAS — Regulation (EU) No 910/2014 — draws a legal distinction that turns out to map unusually well onto the agent-authority question, despite predating it by a decade. Article 26/27 governs electronic *signatures*: instruments that authenticate a natural person's approval. Article 35/36 governs electronic *seals*: instruments that authenticate a legal entity's — an organization's — output, with no natural person's personal approval implied at all.[1] A qualified electronic seal carries a legal presumption of the integrity of the sealed data and the correctness of its origin, attributed to the legal person behind it, and the regulation explicitly contemplates seals supporting fully automated workflows — batch operations with no individual reviewing each instance. That is precisely the shape of the problem an unattended, agent-driven action creates: not "did a person approve this," but "did the organization that authorized this agent stand behind what it just did."

The instrument the industry is building: a way to discover the rules

On 24 June 2026, the American Arbitration Association and Integra Ledger launched the Legal Context Protocol (LCP), with Google, IBM, Circle, Wayfair, and a wide set of other technology and infrastructure companies as founding contributors.[2] LCP solves a different, complementary problem: not "who is legally responsible," but "how does an agent reliably discover and cite the terms, consent basis, and dispute-recourse mechanism that govern a given transaction, before it acts." The base standard is deliberately minimal — a

JSON file published at a predictable path (`/.well-known/legal-context.json`) over plain HTTPS, with the protocol's own documentation stating plainly that no blockchain, cryptographic proof, or centralized intermediary is required at the base tier. That is a notably lighter mechanism than the standard's "trust infrastructure for agentic commerce" framing might suggest, and a genuinely cheap thing for any organization to adopt.

What eIDAS answers

Who is legally speaking when an organization's automated system produces an output with no human clicking "approve." Fully solved for the single-organization case; a decade of settled EU law.

What LCP answers

How an agent finds and cites the actual terms and recourse mechanism governing a transaction before acting on it. A new, lightweight, rapidly-forming standard — not yet widely adopted.

2. What neither one answers

Put the two together and a real gap remains, precisely at the point most "agent-to-agent" framing implicitly assumes is already solved: two *independent* organizations' agents completing a transaction with each other, unattended, in a way that is enforceably binding on both sides. eIDAS tells you who is accountable for one side's automated action. LCP tells an agent where to find the rules. Neither one specifies the handshake — the actual mechanism by which agent A's commitment and agent B's commitment become one mutual, binding agreement, nor how a counterparty is supposed to verify that the agent it's dealing with was authorized to make that specific commitment, within what bounds, before it happened.

In the United States, the doctrinal groundwork for one-sided agent authority is considerably more settled. ESIGN (15 U.S.C. § 7001(h)) and UETA both explicitly recognize "electronic agents" — software acting with no human review at all — as capable of forming enforceable contracts, provided the agent's actions are legally attributable to the party that deployed it. Attribution to the deployer, not prior authorization or bounded authority specifically, is the operative legal requirement; authorization scope and human-in-the-loop thresholds are best practice, not statutory prerequisites.[3] What remains genuinely unsettled — on both sides of the Atlantic — is what happens when an autonomous agent's action turns out to be a mistake: the 2024 *Moffatt v. Air Canada* decision, in which a tribunal held Air Canada bound by an incorrect promise its customer-service chatbot made, is frequently cited as an early, concrete illustration of exactly this exposure.[3] Unilateral-mistake doctrine and UETA's own consumer error-correction provisions offer partial analogies, but neither was written with an autonomous counterparty in mind, and courts have not yet settled how they apply when there is no human on either side of the transaction to catch the error.

THE HONEST SUMMARY

Europe has the stronger single-party instrument. The US has the more settled two-party contract-formation doctrine. Nobody — no vendor, no standards body, no jurisdiction — has yet produced a complete, tested answer for two independent agents forming a mutually binding commitment with each other, verifiably, at machine speed. That is not a gap SignedBy claims to have closed. It is the actual state of the field.

3. A concrete data point from practice

It's worth grounding this in something real rather than purely hypothetical. SignedBy — a European e-signature platform — already operates one small, live piece of infrastructure that sits inside the "solved" half of this picture, and it's worth describing precisely rather than generously. An organization can authorize an MCP-calling agent to certify a document as final: no recipient, no signature request sent to anyone, and no per-document human click, gated by a one-time, organization-level identity verification. It is explicitly not a signature request — it never asks another party to agree to anything; it lets an organization that already considers a document final say so, in an automated, auditable way. That is a real instance of the eIDAS seal model in production, not a thought experiment. Two things are equally important to say plainly: it is a single-organization action, not two agents transacting with each other, and it has not been independently verified to meet the bar for a *qualified* seal specifically, as opposed to a lower advanced or simple tier of assurance — that is a real compliance question this paper does not resolve, and neither has SignedBy, yet. It's offered here as a data point about what the current legal and technical toolkit already makes possible for one side of the transaction, not as a claim about the harder problem this paper is actually about.

4. Open questions worth a real conversation

These are offered as genuine questions, not rhetorical ones. Some of them sit closer to standards-body territory than to any single vendor's roadmap:

1. **The handshake itself.** If eIDAS settles single-party accountability and LCP settles rule-discovery, what's the minimum viable mechanism for two independent agents to form a verifiably mutual, bounded commitment — a bilateral protocol each platform implements separately, or a neutral, shared standard closer to what LCP is already attempting for a related problem?
2. **Mistake correction for autonomous counterparties.** UETA's consumer error-correction provisions and unilateral-mistake doctrine both predate agents acting on both sides of a transaction. Does either framework meaningfully extend to that case, or does it need to be rebuilt?
3. **How much assurance is actually necessary.** Qualified trust service status is a genuine accreditation undertaking, not just an engineering task. For most agent-driven commercial use cases, is that bar the right one, or does an advanced-tier instrument with clear attribution and a solid audit trail already do the job the law actually requires?

4. **Interoperability versus proliferation.** If every e-signature and agent-commerce platform builds its own bilateral trust bridge, that's a lot of duplicated, incompatible plumbing. Is there a real opening for a neutral, multi-party standard — the way LCP already approaches rule-discovery — to cover the handshake too, rather than each vendor solving it alone?

We don't have settled answers to any of these, and we're skeptical of anyone who claims to. We'd genuinely welcome hearing how the people building LCP, and others working on the legal side of agent commerce, are thinking about the handshake question specifically — it's the piece that still feels the most open.

About this paper

Written by the SignedBy team, an eIDAS-oriented e-signature and document platform built in the Netherlands. SignedBy operates live document-sealing infrastructure of the kind described in Section 3, alongside its core e-signature product. This paper reflects independent research and current thinking, not an official position tied to any product roadmap, and nothing in it should be read as legal advice or as a claim of affiliation with the American Arbitration Association, Integra Ledger, or the Legal Context Protocol.

Questions, corrections, or a conversation about any of the open questions above — we'd like to hear them: hello@signedby.ai · signedby.ai

---

1. Regulation (EU) No 910/2014 (eIDAS), Articles 35–36, consolidated text via EUR-Lex: eur-lex.europa.eu/eli/reg/2014/910/2024-10-18/eng.
2. American Arbitration Association, "AAA and Industry Leaders Launch Legal Protocol for Agentic Commerce," 24 June 2026 (adr.org); base standard specification at legalcontextprotocol.org/standard.
3. Chanté Eliaszadeh, "Can an AI Agent Legally Enter Into a Contract?", Astraea Counsel, astraea.law/insights/ai-agent-contract-formation-electronic-agents — including discussion of ESIGN, UETA, and *Moffatt v. Air Canada* (2024).

This paper is provided for informational and discussion purposes only and does not constitute legal advice. Readers with a specific legal question should consult qualified counsel in the relevant jurisdiction.