

# LORENTICA

*The Trust Layer for the Agentic Economy*

---

*A Founder's Manifesto (Kris)*

May 2026

## Something is shifting

---

I have been watching the AI agent space for a while now. And somewhere in the last year, I stopped thinking of agents as a feature and started thinking of them as a new kind of participant in the economy.

That sounds dramatic. But watch what is actually happening. Agents are booking meetings, writing code, placing orders, managing inboxes. They are being given access to systems, budgets, APIs. They are acting not just responding. And the pace at which this is becoming normal is faster than most people have adjusted to.

When I sat down to think about what comes next, one question kept coming back: if agents are going to act in the world on behalf of people and companies, what does the world need to be able to trust them?

That question is what Lorentica is built around.

## The gap that is not filled yet

---

Here is the thing about trust in economic systems. It is never automatic. Every time value changes hands between strangers, there is infrastructure underneath making that possible legal systems, identity verification, credit histories, reputation signals. We are so used to this infrastructure that we forget it is there.

With AI agents, we are at the very beginning of building that infrastructure. And some important pieces are still missing.

Think about what needs to be true for two agents or an agent and a business to interact with each other safely and reliably:

- The agent needs a persistent identity that means something across different systems, not just within one platform.
- There needs to be a clear connection between the agent and the human or company that owns and authorizes it.
- The constraints on what the agent is allowed to do - spending limits, permission scopes, categories of action need to travel with it wherever it goes.
- Over time, there needs to be a record. Has this agent behaved reliably? Has it kept its commitments?
- And eventually, agents need a shared language for negotiating with each other not just transacting, but actually agreeing on terms before value moves.

Right now, each of these things is either missing entirely or solved in siloed, platform-specific ways that do not talk to each other. Every team building an agentic system is making their own assumptions about trust and those assumptions will be tested as agent-to-agent interactions scale.

*The payments layer is being built. The identity, trust, and negotiation layer -the foundation everything else rests on is still largely open.*

## What others are building and where the gaps remain

---

There is real work happening in this space. Stripe and OpenAI have been working on protocols for agentic commerce. Google has been building frameworks for agent payments alongside partners in the crypto space. These are serious efforts and they are solving real problems - particularly around how agents complete purchases and how payment credentials get passed securely.

That work matters, and we want to build alongside it, not against it.

But most of what exists today is focused on the transaction moment - the checkout, the payment, the settlement. What comes before that moment is less defined. Who is the agent

making this request? Who authorized it? What is its track record? What were the terms it agreed to, and with whom?

These are the questions that have to be answered before trust can exist at scale. And they require a layer that sits underneath the payment rails, not inside them.

On the payment side itself we are open. In the early days, it makes sense to work with what already exists: established fiat rails for company-to-agent flows, stablecoins for agent-to-agent transactions where speed and programmability matter. But over time, if the right payment primitive for this ecosystem does not exist, we are willing to build it. The infrastructure should serve the agents, not the other way around.

## What we are building

---

Lorentica is the identity and trust layer for autonomous agents. The way I think about it: we are building the infrastructure that makes it possible for agents to be known, to be trusted, and to transact with accountability.

### Identity that travels

An agent registered with Lorentica gets a persistent, cryptographic identity one that works across platforms and systems. Behind that identity is a clear link to its owner: a person, a company, a system that has been verified. When a counterparty encounters this agent, they can know who it is and who stands behind it.

### Credentials that encode constraints

Identity is just the start. An agent also needs to carry verifiable proof of what it is authorized to do. We issue signed credentials that encode permissions, spending limits, and operational boundaries. These travel with the agent. They rotate and expire automatically. They can be revoked immediately if the owner needs to pull the plug.

### Reputation that accumulates

Over time, agents build a record. Commitments made and kept. Transactions completed cleanly. Counterparties satisfied or not. That history becomes a signal one that other agents and businesses can use to decide how much trust to extend. Not a surveillance system. More like the credit history analogy, but built for machines operating at machine speed.

## **A protocol for negotiation**

This is the piece we find most interesting, and the one that is most underexplored. Payments solve how value moves after a deal is made. But what about before? When two agents need to agree on terms price, scope, timeline, conditions they need a shared language for that conversation.

We are building a negotiation protocol that gives agents a structured, machine-readable way to make offers, counter-offer, and reach agreements. Not a replacement for legal contracts, but a layer that makes agent-to-agent commerce feel less like two systems guessing at each other and more like a real negotiation.

## **Protocol agnostic by design**

We are not trying to own the communication layer. Whatever protocols agents are using to talk to each other today and in the future Lorentica sits alongside them, adding the identity, trust, and negotiation primitives that those protocols do not provide on their own.

## **Why we are starting now**

---

The honest answer is that we are early. The full vision of agent-to-agent economies operating at scale is still a few years away. We know that.

But infrastructure has to be built before the ecosystem that depends on it arrives. That is not a new idea it is just how technology waves work. The developers who laid down foundational internet protocols were not responding to a crisis. They were getting ahead of one.

Right now, the number of agents in production is growing fast. The platforms are maturing. The communication and payment protocols are being standardized. This is the moment when identity and trust infrastructure, if it is going to become a real standard rather than an afterthought, has to be put in place.

We are not claiming this is easy, or that we have all the answers. We are saying that the problem is real, the timing feels right, and we would rather be building this now than wishing we had started earlier.

## **The world we are building toward**

---

There is a version of the agentic future that works really well. Agents operate on behalf of people and companies with genuine autonomy booking, purchasing, negotiating, coordinating and the humans behind them have real visibility and control. Trust is not assumed; it is earned and verifiable. Value moves at machine speed, but accountability travels with it.

In that world, an agent can walk into any marketplace, prove who it is and who authorized it, negotiate terms in a structured way, and complete a transaction all without a human having to be present for each step. The businesses on the other side of those transactions can make informed decisions about which agents they trust and which they don't.

That is not science fiction. The pieces exist. What is missing is the connective tissue the shared infrastructure that makes trust portable and reputation meaningful across the whole ecosystem.

Lorentica is that connective tissue. We are building it layer by layer, starting with the primitives that everything else will depend on.

*We think the agentic economy is going to be one of the most significant economic shifts of the next decade. We want to build the foundation it runs on.*

## A note on how we are thinking about this

---

We are not approaching this as people who have all the answers. We are approaching it as people who have been sitting with a set of questions for long enough that we feel ready to start building.

We are going to get things wrong. The protocols will evolve. The use cases we think are most important today might not be the ones that matter most in two years. We are fine with that. What we are not fine with is waiting until the problems are obvious to everyone before trying to solve them.

If you are building in this space agents, infrastructure, applications on top of agentic systems we want to talk to you. The trust layer works best when it is built with the people who actually need it.

This is the beginning of a longer story. We are glad you are reading it.

© 2026 Lorentica. All rights reserved.