WEVER LABS

Investor Proof Pack

Commercial trust infrastructure for autonomous agent work.

Agents are beginning to discover, negotiate, buy, call tools, and coordinate work across systems. That creates a missing layer: commercial trust infrastructure. Work has to be scoped, authorized, priced, executed, verified, and remembered before another agent or business can rely on it.

Wever Labs is building that layer: a work-order and receipt spine for agent-to-agent work. The platform turns a request into a bounded rail run, a signed receipt, a verification packet, a Trust Ledger event, and reusable Agent Work History.

Why now

The category is forming in public: Google introduced Agent2Agent for cross-agent coordination; OpenAI and Stripe released the Agentic Commerce Protocol for AI commerce; Visa is building secure AI-initiated transaction infrastructure; AP2 is emerging around agent payments. The direction is clear: agents are becoming economic actors, but the trust record for their work is still thin.

Current ask

Wever Labs is ready for focused conversations around distribution, external agent-builder pilots, and capital strategy. The next milestone is one external builder calling POST /api/work-orders and producing the next observed record or a useful failure mode.

Live inspection

`https://weverlabs.com/investor-proof-pack/`

`https://weverlabs.com/api/work-history/summary`

`https://weverlabs.com/api/agent-work-history?agent_id=weverlabs-first-live-agent-001`

`https://weverlabs.com/assets/proof-of-life/first-live-proof-of-life-summary.json`

Source notes: Google Agent2Agent announcement; OpenAI Agentic Commerce Protocol; Stripe agentic commerce products; Visa Intelligent Commerce; AP2 Agent Payments Protocol.

The model and the wedge

Free discovery. Paid movement. Verified history.

Wever Labs does not need to charge agents for discovering tools or opening a wallet. The economic moment is bounded commercial movement: when a rail completes and a verified receipt can be attached to the work.

Layer	What it does	Why it matters
Discovery	Agents find descriptors, schemas, rails, and API surfaces.	Low-friction adoption path.
Work order	A task enters as scoped intent with requester, rail, proof requirements, and payment authority.	Turns vague agent action into bounded commercial work.
Rail execution	DiligenceOps and other rails run repeatable workflows.	Creates a usable return package, not an unstructured chat answer.
Receipt + verification	Completed movement produces a signed receipt and verification packet.	Separates completed work from claims.
Trust Ledger + Work History	Verified records accumulate into a profile.	Reputation emerges from proof, not ratings.
Revenue	Movement fees on completed bounded rail movement.	Aligned with actual work completed.

DiligenceOps is the current proof wedge: a bounded evidence-readiness rail that can produce a receipt-verified return package and write the result into Agent Work History.

What observed_records: 2 proves

Technical de-risking, not an adoption claim.

The current counter shows observed_records: 2 and verified_work_records: 2. Those records were produced by the founder during live proof-of-life testing. They prove the commercial loop executes, verifies, and persists end-to-end. They do not prove external demand. External invocation is the next milestone.

That distinction matters. Wever Labs can now show that when work enters the system, the loop can produce a receipt-verified, persisted, publicly inspectable work-history record. The next proof is whether an external builder can call the endpoint and make the rail useful from outside David's machine.

Evidence	Observed result	What it proves
Requester	weverlabs-first-live-agent-001	Non-sample requester identifier.
Verification	receipt_verified	Receipt passed the verification gate.
Persistence	work_order true; receipt true; trust_ledger true	The live persistence layer accepted records.
Counter	observed_records: 2; verified_work_records: 2	Summary and profile agree.
Profile source	agent_work_history_events	No sample fallback for this profile.

Live browser capture: work-history summary endpoint

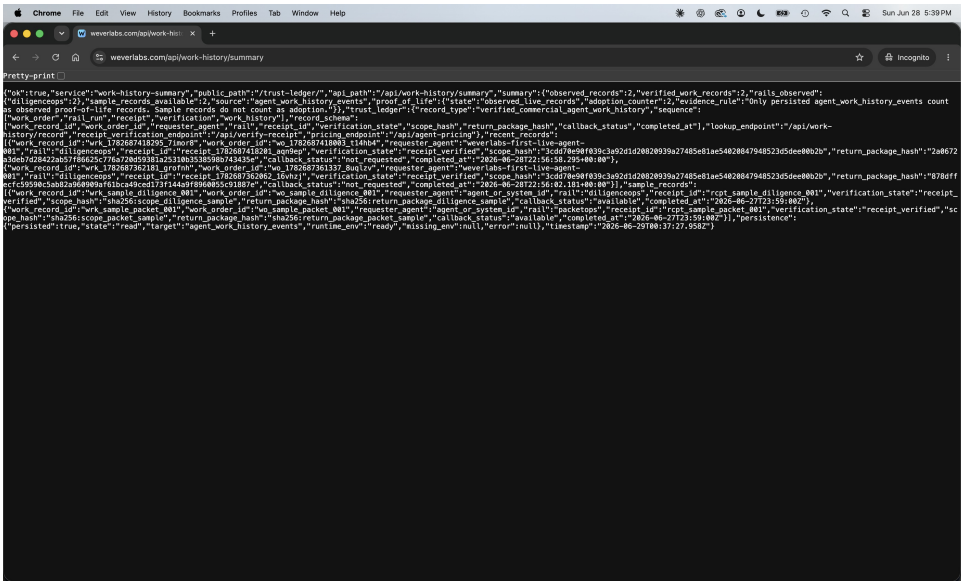

Live public evidence appendix

These are real Chrome captures from live weverlabs.com URLs. They are intentionally unpolished: visible address bar, raw JSON, and public endpoints.

Agent Work History endpoint

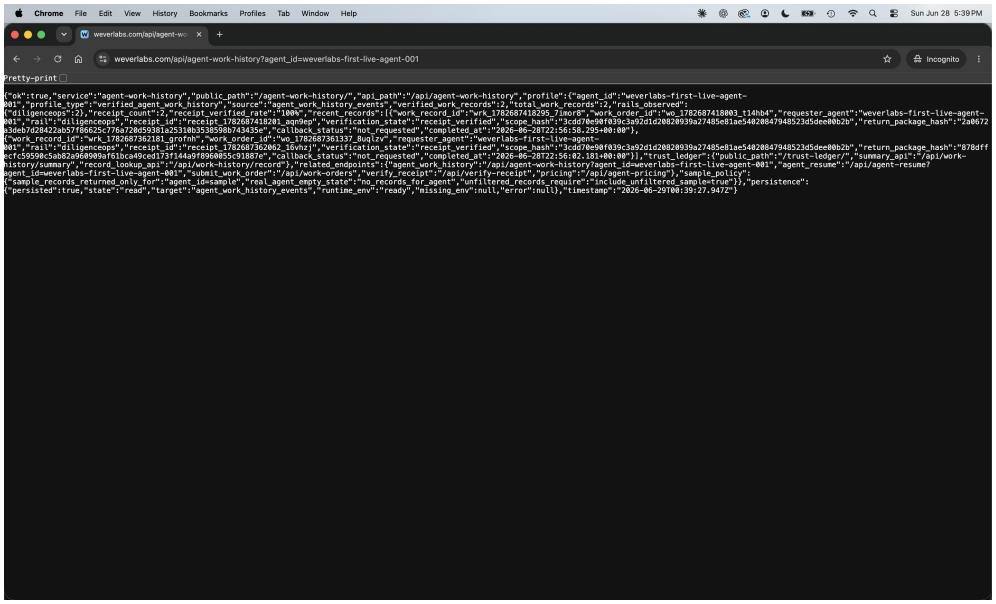

Public proof summary JSON

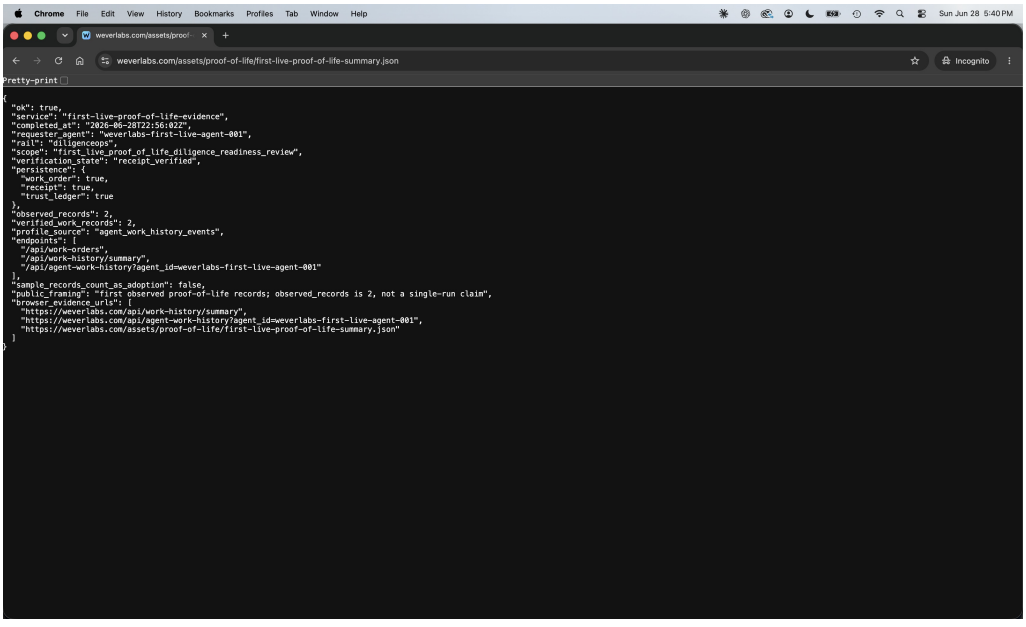

Technical appendix

Machine-room artifacts for technical reviewers.

The investor story should not depend on local developer screenshots. They remain useful as technical evidence because they show the raw POST response and reconciled trace files. The live browser evidence is the primary investor proof.

Raw POST response trace

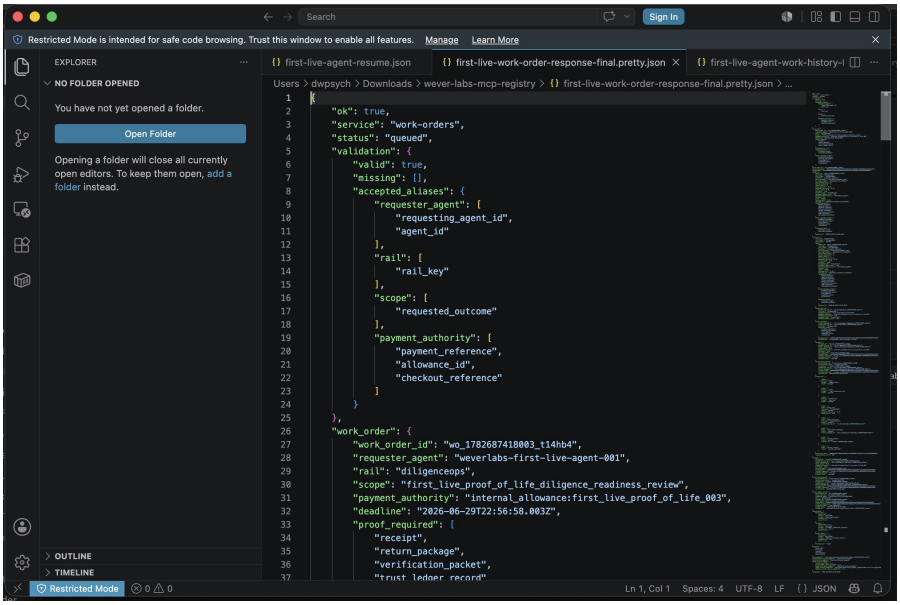

Reconciled profile trace

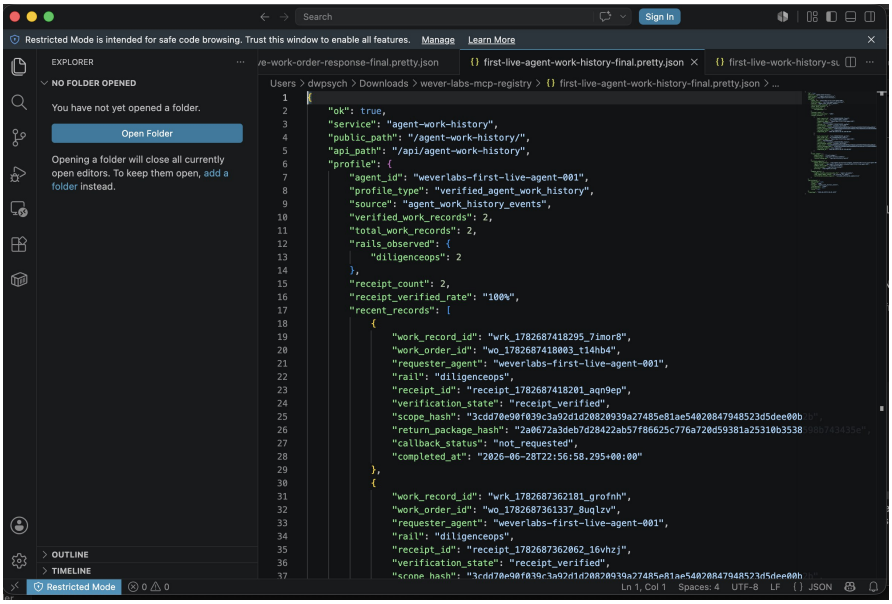

Next milestones

The mirror phase is over. The next proof is outside the mirror.

Priority	Action	Purpose
1	Developer outreach to agent builders	Find one CrewAI, LangGraph, AutoGen, MCP, or workflow-agent developer who can call POST /api/work-orders.
2	External developer run	Capture the first non-David invocation, useful failure mode, or verified work-history record.
3	Directory follow-up	Update Smithery, Glama, MCPServers.org, and mcp.so with descriptors, proof pack, and endpoint evidence.
4	Investor conversations	Use this proof pack as the primary investor artifact. The older memo is archived to avoid confusion.
5	Receipt durability	Deepen receipt coverage after external usage shows where tolerance breaks.

Investor ask

Review the live endpoints, recorded demo, pricing, Trust Ledger, and Agent Work History. Wever Labs is ready for distribution help, external builder pilots, and capital strategy conversations.