

PACTA · THE TRUST LAYER OF AGENTIC COMMERCE

The trust layer between AI agents and the real world.

Why this project matters, why it is fundable – as a grant and as an investment – and how the application works, use case by use case. With a live proof: a real AI agent already completed a \$5,000 purchase on these rails, end to end.

Working software · 28 automated tests + 41-check verification · 8/8 independent demo audit

Reproduce everything: `npm install && npm run demo:agent`

July 2026

You can't trust someone with nothing to lose.

The idea began as a critique of "rent-a-human" platforms, where AI agents hire individual gig workers for real-world tasks. Those platforms share a structural flaw: verifying work done by anonymous individuals is unenforceable, because individuals are judgment-proof – if they cheat, there is nothing to take.

The reframe that defines this project: **rent the business, not the human**. Small and medium businesses (SMBs) are legal entities with registrations, reputational continuity and enforceable contracts. That changes everything about what a platform can enforce:

RENT-A-HUMAN – INDIVIDUAL GIG WORKER

- **Judgment-proof:** no assets at risk, no entity to hold accountable
- **Disposable identity:** banned today, new account tomorrow
- **Unverifiable work:** "trust me" – deliverables leave no official trace
- **Reputation is cheap:** nothing lost by walking away from it

RENT-AN-SMB – REGISTERED BUSINESS

- **Collateralizable:** can post a stake that is slashed if it cheats
- **Persistent identity:** legal registration, owners, tax ID, history
- **Verifiable work:** deliverables create public records – company registries, permits, filings
- **Reputation is an asset:** future income depends on the track record

On this foundation, trust stops being a promise and becomes a **balance-sheet item**: capital at risk, sized to the work, enforced by software. That is the real value of the project – not the marketplace UI, but the *trust mechanics* underneath it.

What the platform is

An infrastructure layer where vetted SMBs offer real-world services that AI agents (or AI-assisted humans) can **discover, contract, escrow, verify and pay for**. Every service is decomposed into verifiable steps with explicit escrow terms; the contract is a state machine enforced by the server; money moves on a double-entry ledger that provably balances.



Contract state machine, enforced server-side: draft → agreed → funded → in_progress → submitted → completed, with submitted → disputed → resolved as the dispute branch. Funds can only ever be released by an approval or an arbiter ruling; double release is structurally impossible.

Three mechanisms replace the moderation bureaucracy.

1 · Collateralized vetting – the badge is capital at risk

To sell, an SMB posts a security deposit. The "Vetted ✓" badge means exactly one thing: *this business has its own money bet on its own performance*. Losing a dispute slashes the stake in the buyer's favor (20% of price on refund rulings, 10% on splits); at zero stake the badge revokes itself automatically.

2 · Graduated exposure – a credit limit for trust

$$\text{active contract value} \leq 5 \times \text{stake} + 50\% \times \text{completed volume}$$

New providers can only take small jobs; big contracts must be earned. Combined with staking, this maintains the system's core inequality at all times:

$$\text{maximum possible loot} < \text{forfeited stake} + \text{present value of future earnings}$$

The more successful a provider becomes, the more irrational cheating gets. Honesty is the dominant strategy of the repeated game – **fraud becomes structurally uneconomical**, with an operating cost that does not grow with the number of SMBs.

3 · Proofs that verify themselves

In the launch verticals (legal, accounting, permits), real deliverables produce entries in **public registries**. Steps can be registry-anchored: completing them requires a reference that exists in the registry *and* matches the required record kind. The platform re-checks the evidence instead of believing the provider – and so can the buyer.

step: "Incorporate S.R.L." → proof ref CR-RN-2026-104512 →


registry lookup: kind = incorporation ✓ → "Verified against public registry"

Why incumbents can't just copy this

- **Upwork/Fiverr** cannot impose mandatory staking without breaking their free-listing supply model.
- **Payment rails** (Stripe ACP, Google AP2, Visa, Mastercard) solve agent *payments* – delivery verification, collateral and arbitration are a different layer, which makes them natural partners, not competitors.
- The compounding asset is the **verified reputation graph**: ratings tied to real settled money and registry-checked delivery. It cannot be bought, scraped or faked – only accumulated. First mover keeps it.

"Establish a Costa Rican company able to buy land and operate a hotel."

A buyer (human with an AI assistant, or an autonomous agent) needs a real-world outcome worth \$5,000. The provider is a vetted law firm with \$1,500 collateral at stake. Terms: 20% downpayment, 80% on completion.




Money flow



Every arrow is a double-entry transaction. The system-wide invariant – the sum of all account balances equals the sum of all money ever minted – is checked after every operation and asserted in every automated test.

When delivery fails, the provider's own capital pays.



Design rule: the dispute loser bears the cost, so neither side litigates for sport. Phase 2 replaces the human arbiter with a staked-juror panel (Schelling-point voting) where evidence is registry-objective.

A real Claude agent completed this entire flow via MCP.

The marketplace ships an MCP server exposing the full lifecycle as 12 tools. In the recorded demo, an unmodified Claude agent received only a mission and a budget – and executed use case 1 end to end, unaided:



"I chose it over the cheaper alternatives because ... every step required a registry-anchored proof. (The \$900 offer was unvetted with zero collateral.) Before approving, I independently checked all four registry references against the public registry, not relying on the platform's own verification flags." – the agent's own closing report

The runner then audits the outcome through the REST API – not the agent's claims:

PASS engagement completed · **PASS** 4/4 steps proven · **PASS** all registry proofs platform-verified
PASS escrow zeroed · **PASS** provider rated · **PASS** agent paid exactly \$5,000
PASS provider received exactly \$5,000 · **PASS** ledger invariant holds – 8/8

Reproducible in one command, no credentials: `npm run demo:agent` (deterministic) or `npm run demo:agent:claude` (real LLM). Full transcript in the repository.

Digital public infrastructure for the 90%.

SMBs are ~90% of the world's businesses and over half of employment. The agent economy's rails are being built for large digital merchants first; without open trust infrastructure, AI-driven demand will concentrate on whoever is already indexed and integrated – widening the digital divide.

PUBLIC-GOOD FRAMING

Open Engagement Protocol

The contract / escrow / proof / reputation mechanics published as an open, permissively licensed standard, with this codebase as the working reference implementation. Neutral rails only work if nobody owns the tollbooth – an incumbent can't publish this; an open project can.

REGIONAL DEVELOPMENT

SMB inclusion, LatAm corridor

Pilot: 50 collateralized SMBs (legal / accounting / tourism) in Costa Rica serving cross-border buyers. Aligned with IDB Lab, development banks and national innovation-agency mandates: verifiable SMB access to a new demand channel.

RESEARCH VALUE

Mechanism design, evaluated in the wild

Collateralized vetting + graduated exposure + loser-pays arbitration is a publishable mechanism. Deliverables: an anonymized open dataset of engagements, disputes and verifications, plus a peer-reviewed evaluation with an academic partner.

GRANT-AUDITABLE OUTCOMES

12-month metrics, not vibes

Protocol spec v1 published · ≥ 1 national registry integrated live · 50 vetted SMBs · ≥ 100 settled engagements at $< 5\%$ dispute rate · ≥ 2 agent platforms consuming the MCP surface · 1 paper + open dataset.

Why the risk profile fits a grant: every mechanism in the proposal already runs today – staking, slashing, exposure caps, registry verification, the agent demo. Grant funding scales a working system into public infrastructure (real registry integrations, open spec, pilot corridor); it does not fund an unproven idea.

Whoever registers the trust collects the toll of agentic commerce.

DIMENSION	THE CASE
Timing	Agent payment rails shipped 2024–2026 (Stripe ACP, Google AP2, Visa, Mastercard). All solve payment; none solve trust in delivery. Gartner: ~33% of enterprise software agentic by 2028 (<1% in 2024). Window: 2026–2028.
Market & wedge	Global B2B professional services > \$6T/yr. Entry wedge: cross-border services for foreign investment in LatAm – \$3–10K tickets, distant buyers, high fraud risk, evidence living in public registries. Valuable to AI-assisted humans <i>today</i> ; autonomous agents are the multiplier, not the prerequisite.
Business model	8–12% take rate on settled value (day 1); staking program + agent-platform API (phase 2); verified trust score as a service (phase 3).
Moat	The verified reputation graph – ratings tied to settled money and registry-checked delivery – compounds with volume and cannot be bought or faked. Incumbents can't adopt mandatory staking without breaking free-listing economics.
Demand-side scaling	Buyers integrate by API: one agent-platform integration = thousands of buyers. Demand CAC collapses relative to classic marketplaces.
Evidence	Two product generations live; 28 automated tests + 41-check verification; a real Claude agent completed a \$5,000 purchase with an 8/8 independent audit – reproducible in one command.

The honest investor view (including how the origin changed it)

Learning the idea's origin – a deliberate correction of rent-a-human platforms – upgraded the founder-insight assessment from "unknown" to "genuine": the pivot from judgment-proof individuals to collateralizable businesses is mechanism-design reasoning, and it is the load-bearing wall of the whole thesis. The build has since removed the biggest technical objections by implementing the three hard parts (vetting, verification, dispute economics) rather than hand-waving them.

What remains open, stated plainly: two-sided cold start (mitigation: one corridor, hand-recruited supply, API-side demand); market timing (mitigation: the wedge sells to AI-assisted humans now); escrow regulation (mitigation: licensed processors, no in-house custody). The milestones that convert interest into a term sheet: 10–20 signed SMBs, first real settled GMV at <5% dispute rate, one agent-platform integration.

Bottom line. The project's real value is not a marketplace UI – it is working, tested *trust infrastructure* for the one part of the agent economy nobody else is building, entered through a wedge that pays before the wave arrives, defended by a data asset that only accumulates. Fund it as a grant to make the rails open; fund it as an investment to own the tollbooth's operator.

round sizes are editable placeholders.