REALITY AS A SERVICE™

PROXY

The physical-world data layer for AI agents

01

AI can't touch reality. We fix that.

02

Every answer returns as structured, verifiable truth.

03

We're building the data layer for the physical world.

useproxy.io

THE PROBLEM

AI can reason about anything.

It can't touch the world.

Every agentic workflow eventually hits the same wall. The moment an agent needs to verify something in the physical world — a drug in stock, a price on a shelf, a condition at a site — every existing tool fails.

No database. No API. No model can substitute for a human physically present at a location at a known time with provenance.

AI can reason

Plan, analyze, decide with remarkable precision across any domain.

AI cannot observe reality

No model can walk into a pharmacy, check a shelf, or verify a physical condition.

✅ **Proxy closes that gap**

Human sensors act as the physical observation layer — returning structured, verified truth.

INDEPENDENT VALIDATION

Three voices.
Same conclusion.

> *"I'm surprised we don't have information markets where taking a photo from somewhere costs $10. Agents trying to resolve real-world positions need that."*
>
> **Andrej Karpathy**
> No Priors Podcast · March 2026

> *"Agentic scaling — the fourth scaling law — depends on agents accessing ground truth to perform research. Physical ground truth is the missing layer."*
>
> **Jensen Huang · NVIDIA CEO**
> GTC 2026 · Lex Fridman Interview

> *"Payment is solved. Discovery is being solved. Physical verification — did it ship, is it on the shelf, was the condition as claimed — that layer doesn't exist yet."*
>
> **Sam Ragsdale**
> Open Agentic Commerce · 406K Views

Three independent conversations. Same conclusion. The category is forming now.

THE INEVITABILITY ARGUMENT

This layer will exist.

The only question is who owns it.

01

Agents will need to verify — not eventually, structurally

Physical verification is not a feature request. It is a constraint of reality. Every agent that acts in the physical world will eventually hit the same wall. That wall does not go away as models improve.

02

The cost of wrong data is quantifiable and immediate

A wrongly-resolved market position has a dollar loss with a timestamp. A misrouted patient has a measurable delay. A failed compliance audit has a remediation cost. The ROI calculation is never ambiguous.

03

No existing system can substitute for physical observation

Search engines, data brokers, web scrapers — none of these return what was physically observed at a specific location at a specific time with a chain of custody. The gap is architectural, not a data gap.

04

The data moat compounds only for the first mover

The Ground Truth Database becomes exponentially more valuable with each query. The company that owns the data owns the category. That window does not stay open.

This is not a product bet. This is an infrastructure bet.

THE SOLUTION

Proxy
Reality as a Service™

When an agent needs to know something about the physical world, Proxy dispatches a vetted human sensor, collects structured proof, and returns confidence-scored, provenance-backed ground truth — in minutes.

01

Query

Any agent, operator, or autonomous system posts a structured query with location, query type, and price. Agents can query, set the price, make the payment, and manage verifications autonomously — no human in the loop.

02

Dispatch

Geo-matching engine finds the closest verified Proxy Sensor. Push notification sent within 60 seconds. Stripe escrow holds payment on acceptance.

03

Truth Returned

Structured JSON with confidence score, cryptographic provenance record, human-readable summary, and Truth Atoms stored permanently in the Ground Truth Database.

ONE REAL QUERY

Five steps.
One verified answer.

01

Agent asks

"Is Ozempic 0.5mg in stock within 5 miles of Salt Lake City?"

02

Proxy dispatches

Geo-match → sensor 0.4mi away → push notification → escrow held in 8 seconds

03

Human observes

📍 CVS, 400 N 900 E

📷 Geotagged photo proof submitted

04

Truth returned

{ "in_stock": true, "confidence": 0.91, "resolved_in": 18 }

05

Decision made

"Confirmed. Routing patient to CVS, 400 N 900 E, SLC."

✅ Ozempic in stock at CVS. Verified 18 min ago. Confidence 0.91. Patient routed. · 6 Truth Atoms stored permanently.

THE MOAT

The Ground Truth Database *compounds with every query.*

The same pattern Jensen built with CUDA: maximize the install base, build the ecosystem, compound with use. Every Truth Atom stored makes the next query more accurate. The database IS the moat.

Truth Atoms

Every query decomposes into typed, decaying, provenance-hashed atomic units stored permanently. The database compounds with every observation.

Confidence Scoring

Accuracy vectors per sensor per query type, updated by outcome feedback. Not star ratings — dynamic priors that improve with use.

Truth Decay Model

Exponential decay with calibrated half-lives by query type. Physical truth ages — Proxy tracks it and alerts before stale data drives a decision.

Shadow Queries

5% of dispatches are known-answer tests. Sensors never know which queries are tests. The deterrent is uncertainty — gaming the system is irrational.

Cryptographic Provenance

SHA-256 hashed chain of custody from observation to output. Evidence is presentable, not just referenceable — critical for compliance and audit.

Human Translation Layer

Every JSON output also returns a human-readable sentence for operators. Two outputs, always — agents read the JSON, humans read the sentence.

THE ANCHOR WEDGE

Why information markets *are the opening move.*

01 **Every position resolves to a physical fact**

Polymarket and Kalshi markets don't resolve on opinions. They resolve on observable physical reality — a building completed, a product launched, an election decided. Agents can aggregate predictions with precision. They cannot independently verify the fact that closes the position.

02 **The cost of wrong data is immediate and quantifiable**

A wrongly-resolved position has a direct dollar consequence. Unlike most data problems where staleness is inconvenient, here it's a financial loss with a timestamp. That makes the verification fee instantly justifiable — the ROI calculation is trivial.

03 **Karpathy named this gap unprompted**

On the No Priors podcast, Karpathy described exactly this use case without knowing Proxy existed. That's not a trend. That's a structural gap in the agent stack that the most credible voice in AI identified independently.

Markets that need Proxy

Polymarket
Physical event resolution

Kalshi
Real-world outcome verification

AI Trading Agents
Ground truth for position logic

Prediction APIs
Autonomous market settlement

Sports / Election Books
Physical outcome confirmation

USE CASES

Built for the verticals *where truth is expensive.*

INFORMATION MARKETS

Physical verification for market resolution & AI trading

Prediction markets, AI trading bots, and automated settlement systems aggregate opinion with precision. Their failure point: they cannot verify the physical-world fact that resolves the position.

› Verify construction milestone at this address
› Confirm product launch availability at retail
› Physical event attendance at venue X

COMPLIANCE & AUDIT

SOC-1, site audits & regulatory verification

Compliance evidence must reflect what was physically observed at a specific location at a specific time. Proxy returns cryptographically chained evidence from sensor observation to structured output — presentable, not just referenceable.

› Confirm physical security controls at facility
› Verify storage conditions at dispensing location
› Confirm site readiness for regulatory inspection

RETAIL & INVENTORY

Shelf-level commercial intelligence at scale

Brands spend billions assuming their retail execution matches their plans. It rarely does. Proxy closes the gap between what a brand believes is happening at the shelf and what a sensor confirms is actually there.

› Is Product X actually on shelf at this location?
› Verified shelf price vs. advertised price?
› Is the promotional display executing as planned?

COMPETITIVE LANDSCAPE

They are task platforms.
Proxy is data infrastructure.

	Proxy	RentAHuman	DoorDash Tasks	MTurk / TaskRabbit	IoT / Cameras
Structured JSON output	✅	❌	❌	❌	⚠️ fixed schema only
Confidence scoring	✅	❌	❌	❌	❌
Cryptographic provenance	✅	❌	❌	❌	❌
Truth Decay model	✅	❌	❌	❌	❌
Agent-native API / MCP	✅	❌	❌	❌	❌
Compounding data moat	✅	❌	❌	❌	⚠️ siloed data
Shadow query trust system	✅	❌	❌	⚠️ partial	❌
Instant global deployment	✅	⚠️ limited	✅	✅	❌ capex required

The question is not 'why not DoorDash + Claude?' — it's why no one else stores the structured, confidence-scored, provenance-hashed output permanently. That database is Proxy's CUDA.

THE SUPPLY ENGINE

Sensors recruit themselves.
The receipt is the ad.

HOW SUPPLY IS ACQUIRED

PATH 1 — **Receipt posts become worker ads**

Every X post showing '$40 earned in 18 minutes' targets the exact people who want to earn between DoorDash runs. The content engine and the worker acquisition engine are the same thing.

PATH 2 — **Existing gig workers run Proxy independently**

DoorDash, Uber, and Instacart drivers are already positioned across the city. Proxy runs as a fully independent app alongside existing gig platforms — their routes are already right.

QUALITY CONTROL MECHANISMS

Shadow Queries
5% of dispatches are known-answer tests. Sensors never know which.

Accuracy Vectors
Per query-type historical accuracy score updates after every submission.

GPS Anti-Spoofing
Velocity detection and device fingerprinting catch location fraud.

Photo AI Check
Claude reviews every photo submission for schema consistency.

Scale path: SLC — 4 sensors live → Santiago, Chile — next → Gig worker integrations → 10+ cities by end of year

TRACTION

The loop is running.

These are the receipts.

PILOT LIVE — SALT LAKE CITY, UT

Seeding queries running · 0% failure rate · 4 sensors, 100% dispatch acceptance rate to date

4	0.91	18 min	SLC → SCL
Sensors registered	Avg confidence score	Avg resolution time	Pilot cities

`market_resolution · Denver, CO`

Agent asked: "Verify construction milestone completion at 1420 Larimer St, Denver — market resol on physical confirmation."

Proxy dispatched a sensor to the site. Sensor submitted geotagged photo proof of completed foundation pour with structural framing visible. Result returned in 28 minutes.

✅ Milestone confirmed. Verified 28 min ago. Confidence 0.94. Market resolves YES.

0.94	confidence
28 min	resolution
4	truth atoms

THE ROI ARGUMENT

The cost of NOT verifying *always exceeds the cost of verifying.*

	Without Proxy	With Proxy	
Information Market	$10,000 position resolved incorrectly = $10,000 loss, timestamp-stamped	$40–120 verification query	**99%+ cost avoidance**
Compliance & Audit	SOC-2 or regulatory finding from unverified site condition = $50,000–$200,000 remediation + reputational cost	$60–200 per site verification	**99%+ cost avoidance**
Retail & Brand	Undetected promotional non-compliance across 500 stores = $2M–$5M in lost promotional ROI per campaign	$15–40 per store verification	**500x–1000x ROI**

Agents optimize for certainty. When the cost of uncertainty is quantifiable, Proxy becomes the obvious spend.

BUSINESS MODEL

Infrastructure pricing.
Every layer compounds.

Standard Query Per-query pricing. Agent sets price at submission. Higher price = faster dispatch priority.	$15–40 / query	15%
Information Market Verification Dedicated market resolution queries: milestone confirmation, product launch verification, physical event attendance.	$40–120 / query	15%
Reality Monitor Subscription Ongoing monitoring with decay-triggered alerts. Agents subscribe to physical reality for key entities.	$99–999 / mo	Full margin
Ground Truth Data License Enterprise access to the compounding Ground Truth Database. The moat monetized directly.	$500–10K / mo	Full margin
Enterprise API Dedicated vertical API for prediction market platforms, compliance teams, and enterprise agent operators.	$2K–25K / mo	Full margin

Path to $1M ARR: 500 queries/day at $30 avg = $820K/year + subscriptions + data licenses.

WHY WE WIN

Category dominance *is decided in the first 18 months.*

01

Domain authority that cannot be manufactured

Joshua operated PURE Healthcare — a specialty infusion company he co-founded and scaled before merging with AleraCare. He watched the physical verification gap cost time, money, and patient outcomes in real workflows for years before building Proxy. That's not a whiteboard thesis. That's a founder who lived the problem in the most demanding vertical possible.

No competitor without that operational background can manufacture equivalent credibility with healthcare or compliance buyers.

02

First data advantage that cannot be purchased

The Ground Truth Database compounds from Day 1. Every query makes the confidence scores more accurate. Every sensor's accuracy vector improves. Every decay curve gets calibrated. A competitor who starts today doesn't just lack the technology — they lack the data. That gap widens with every query we resolve.

This is the same moat Jensen described with CUDA: the install base defines the architecture. You can't buy your way in after the fact.

03

Speed of proof that signals urgency

Chase Dobson built the complete Proxy stack — backend API, PostGIS geo-matching, React Native sensor app, Stripe Connect payments, confidence scoring, Truth Atom storage, shadow query infrastructure — in under one calendar week. The loop is already closing.

We are not pitching a concept. We are demonstrating a working system at the moment the category is being publicly named by Karpathy, Jensen, and Ragsdale simultaneously.

RISKS & MITIGATIONS

Every risk is architectural, *not operational.*

DATA QUALITY — Wrong data leads to bad decisions

Shadow queries: 5% of dispatches are known-answer tests. Sensors never know which — uncertainty is the deterrent.

Accuracy vectors: Per-sensor, per-type scores update after every outcome. Bad sensors filtered automatically.

AI photo check: Claude reviews every submission for schema consistency before result is returned.

Confidence scoring: Every result carries scored probability — 0.91, not just 'verified'. Low confidence triggers re-query.

LIABILITY — Wrong output causes measurable harm

Observational data only: Workers collect publicly observable retail facts. J.D. framework confirms this is not HIPAA-regulated.

Cryptographic provenance: SHA-256 hashed chain of custody. Immutable record shifts liability appropriately.

Per-query disclaimer: Every instruction states: 'This query requests only publicly observable information.'

SUPPLY QUALITY AT SCALE — Sensor reliability degrades as network grows

Tiered access: High-accuracy sensors get priority dispatch and better pay. System self-selects for quality.

GPS anti-spoofing: Velocity detection and device fingerprinting catch location fraud before it reaches output.

Worker isolation: Sensors never know which queries are shadow tests. Gaming the system is structurally irrational.

REGULATORY / LABOR — Worker classification or data privacy challenges

Independent contractor model: Sensors complete discrete observation tasks — same structure as gig platforms globally.

Entity separation: Proxy is legally distinct from all other ventures. Confirmed by founder counsel (J.D.).

Platform ToS: Standard SaaS framework covering worker data, query outputs, and retention — ready before launch.

Every failure mode has a corresponding architectural solution. Risk is not ignored — it is designed around.

THE TEAM

Built by an operator *who lived the problem.*

Joshua Tandy, J.D.

Co-founder & Managing Partner
BUILD Capital Partners

Federal court clerk, Judge David Sam, U.S. District Court, Utah. Goldman Sachs private wealth management. Co-founder PURE Healthcare, merged with AleraCare. Co-founder and Director, Memory Treatment Centers (CerebraLinked Platform for AD); Board Designee Aura TopCo (PE-backed). Built AI trading systems on Polymarket.

- Clerk — Judge David Sam, D. Utah
- Goldman Sachs — Private Wealth
- Co-founder, PURE Healthcare / AleraCare
- Co-founder & Director, Memory Treatment Centers
- Board, Aura TopCo (PE-backed) · @jtandylaw

Chase Dobson

Lead Developer — AI Intern
BUILD Capital Partners (Jan 2026)

Recently returned from a two-year mission in Panama — fluent Spanish and Dulegaya. Built the entire Proxy stack in under one calendar week: backend API, PostGIS geo-matching, React Native sensor app, Socket.io sync, Stripe Connect escrow, confidence scoring engine, Truth Atom storage, and shadow query infrastructure.

- Mission — Panama · Spanish & Dulegaya
- AI Intern, BUILD Capital Partners (Jan 2026)
- Full-stack: backend, mobile, web portal
- PostGIS dispatch · Stripe Connect · Truth Atoms
- 30-day build plan in ~1 calendar week

AGENT ECONOMICS

Agents hold budgets.
They optimize for certainty.

The economic model that makes Proxy inevitable: autonomous agents are being given budgets and the authority to spend them. An agent optimizing for correct outcomes will always choose to verify over guess — because the cost of a wrong decision exceeds the cost of verification by orders of magnitude. Proxy becomes a default spend primitive in every agentic stack.

x402 + mpp

Payment protocols for agents

The payment layer is already solved

Coinbase's x402 and Stripe's mpp protocols give agents a permissionless way to spend money. The infrastructure for agents buying services exists now. Proxy is the physical verification service they will buy.

$0.00

Human intervention required

Fully autonomous query lifecycle

An agent can discover the Proxy API, post a query, set the price, hold escrow, receive results, and submit outcome feedback — with zero human in the loop. Every step is autonomous by design.

↑ EV

Expected value of verification

Agents are rational buyers

Agents evaluate expected value. When the cost of wrong data is quantifiable — a $10K position, a $200K regulatory finding — and the verification cost is $40–120, the math is never ambiguous. Proxy becomes the rational default.

THE DEFENSIVE ARGUMENT

Won't robots and cameras *eventually replace this?*

No. Here is the precise argument.

Humans deploy instantly. Robots don't.

A camera network or robot fleet requires capital, installation time, permitting, maintenance, and geographic pre-commitment. Proxy sensors are people who already live in every city on earth. Geographic coverage is instantaneous. A robot cannot be in Salt Lake City tomorrow. A Proxy sensor can be dispatched in 60 seconds.

Humans handle edge cases. Machines break on them.

Physical verification often requires judgment — a label that's partially obscured, a stairwell that's oddly lit, a situation that requires asking a staff member. IoT sensors return binary states for pre-defined conditions. Proxy sensors return structured truth for whatever the agent actually needs to know.

The data moat is sensor-agnostic.

Even in a world where IoT density increases, the Ground Truth Database compounds. Every observation — human or eventually machine — stores Truth Atoms. If IoT sensors become cheap and ubiquitous, Proxy integrates them as a verification source. The architecture is designed for this. The moat is the data, not the human specifically.

MARKET SIZE

Every agent that interacts *with the physical world.*

The addressable market is not defined by current query volume — it is defined by the number of agentic workflows that will eventually require physical verification. That number is every agent that acts in the physical world. The question is not the size of the market. It is the speed of capture.

$2B+	**SAM — Information Markets** Prediction market verification, AI trading agents, autonomous settlement systems. Physical resolution of market positions is structurally required. Clear, immediate, quantifiable value per query.	Beachhead. First 18 months.
$15B+	**SAM — Enterprise Verticals** Healthcare verification, compliance and audit, retail intelligence, logistics confirmation. High value per query ($40–$200). Institutional buyers with clear compliance requirements.	Phase 2. Post-pilot.
$100B+	**TAM — Physical Intelligence Layer** Every agentic workflow that needs to verify anything in the physical world. As agent adoption scales — Jensen's fourth scaling law — the demand for physical ground truth scales with it.	Long-term category.

Market sizing premise: Jensen is right about agentic scaling. Karpathy is right about physical verification demand. All three validators identified this independently.

THE FLYWHEEL

Every city gets better *as density compounds.*

City Bootstrap Logic

- **SLC — 4 sensors**
  Pilot proves loop
- **10 sensors**
  Avg dispatch < 10 min
- **25 sensors**
  Full city coverage
- **50+ sensors**
  Sub-5 min SLA live
- **Santiago next**
  Gonzalo leads supply

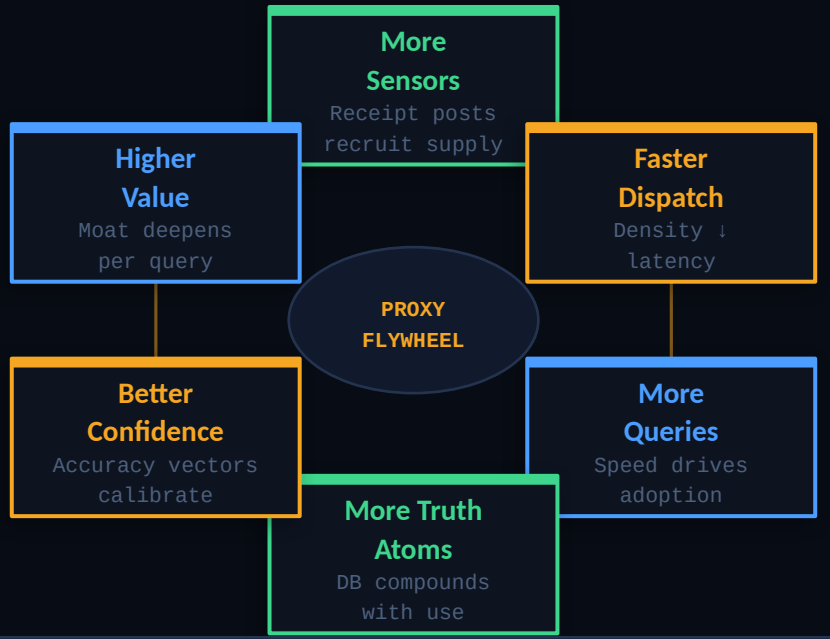

The flywheel is not theoretical. It is running in SLC today. Every query tightens the loop.

The category clock is running.

Every agent will eventually need to *query reality.*

We intend to be the layer they query.

Agent Builders — LangChain, CrewAI, Claude ecosystem developers who need the physical verification layer.

Information Market Operators — Polymarket, Kalshi, and AI trading platforms that need independent ground truth for position resolution.

Strategic Partners — DoorDash, Uber, Instacart ecosystem — worker supply infrastructure for geographic expansion.

APPENDIX

STRATEGIC CONNECTIONS

Not raising. *Connecting.*

Bootstrap through Day 30. The demo receipt is worth more than a deck right now. We are not raising before the loop is proven on record. When we do raise, it will be on our terms — after the flywheel is visibly turning.

AGENT BUILDERS

LangChain, CrewAI, Claude ecosystem, MCP server developers

API integration partners who need the physical verification layer. First developers to integrate Proxy into their stack become design partners who shape the product roadmap.

Ask: Early API access. Feedback on the query schema and MCP integration.

ENTERPRISE OPERATORS

Prediction market platforms (Polymarket, Kalshi), compliance teams, healthcare operators

First enterprise customers who pay for query access. A single named reference from this category transforms how every subsequent conversation goes.

Ask: Pilot query access. Letter of intent or informal commitment to test.

STRATEGIC PARTNERS

DoorDash, Uber, Instacart ecosystem — approached post-loop-proven

Worker supply infrastructure for geographic expansion. Not a tech partnership — a data licensing relationship. They pay for enterprise API access to query the Ground Truth Database per city.

Ask: Introduction to operations or data teams. Andy Fang (DoorDash) is a warm door.

APPENDIX

BUILD ROADMAP

The loop closes.
Then the product ships.

~1 day	to demo
~1 week	operator portal
~2 weeks	full platform

Phase 0	**Infrastructure complete** Backend API · PostGIS · Stripe Connect · React Native sensor app · Socket.io · Shadow queries · Truth Atoms	Done ✓
Phase 1	**Day 30 demo — loop on record** Claude agent → CVS pharmacy → JSON output → routing decision → Stripe payout → posted to X	~1 cal. day
Phase 2	**Operator web product (build first)** Public signup · query submission UI · output viewer · API key management · billing · first 10 paying customers	3–5 cal. days
Phase 3	**Sensor app polish** Public onboarding · earnings dashboard · receipt sharing to X · only when volume exceeds 4 sensors	2–3 cal. days
Phase 4	**Full two-sided platform** Self-serve · autonomous agent integration (x402/mpp) · MCP server · Reality Monitor · Santiago launch	6–10 cal. days
Phase 5	**Reality Graph + enterprise scale** SDK · dynamic routing · swarm mode · Decision Attribution Layer · multi-vertical expansion	Post PMF

Calendar days reflect Chase's demonstrated pace (8–9 plan days per calendar day). Full 8-slide roadmap available separately.