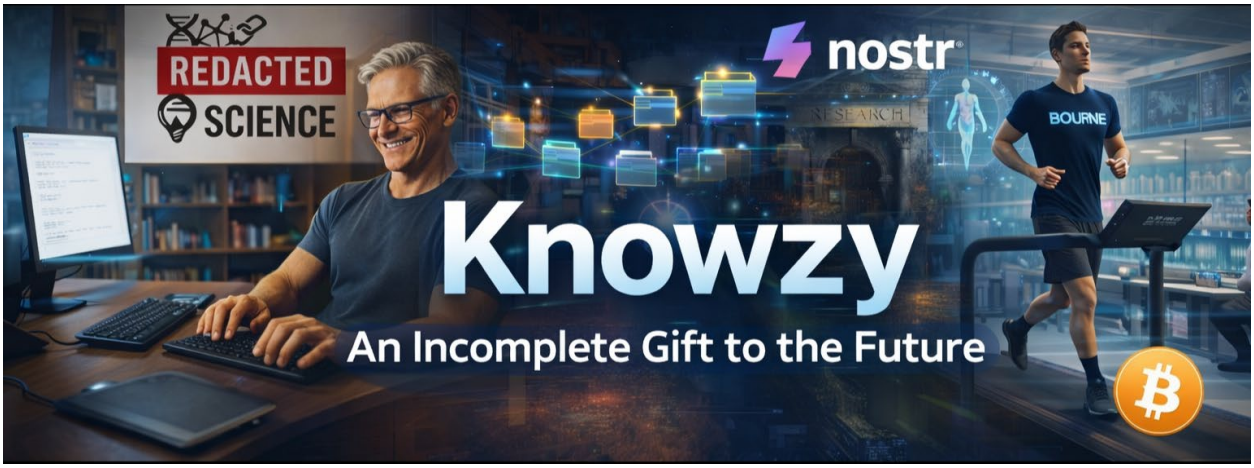

Homepage: jimcraddock.com Home of Redacted Science the Story of Coevolution

Seven years ago, I designed a polling platform where respondents own their data and earn sats for their answers. Life intervened, but with nostr and Lightning, the concept is viable again. Knowzy flips traditional surveys: you build a dimensional profile through your answers, question askers target specific audiences transparently, and payment flows to those who respond. Consistency tracking over time creates a marketplace for coherent perspectives. Your profile is portable, your reputation is earned, your data is yours. The architecture is sound - flexible dimensional modeling, Lightning payments, nostr identity. The database schema is complete. What's missing is implementation: the frontend, the Lightning integration, the seed questions that bootstrap new users. I'm sharing this knowing I probably won't finish it. But incomplete work has value if it's documented. The repo is public at https://github.com/rjcraddock/knowzy The data model is there. The vision is clear. For whoever comes next: he who controls the indexes controls the future. This is an attempt to build an index users actually control. Build it better than I could have.

Knowzy: An Incomplete Gift to the Future

I've been thinking about indexes for a long time. Decades, actually. In my book, I wrote that "he who controls the indexes controls the future" - and I meant it literally. The way we organize, categorize, and retrieve information shapes everything that comes after. Every recommendation you see, every search result, every "people like you also..." is an index someone else built, optimized for their interests, not yours.

Seven years ago, I started designing Knowzy - a polling platform where the respondents own their data, the indexes are transparent, and the incentives align with truth rather than engagement. Life got in the way. Health issues I've documented elsewhere derailed the project. But in 2026, with AI-assisted development and nostr providing the infrastructure I would have had to build from scratch, the concept is viable again.

If you've read my book, Redacted Science, you understand my passions: truth, architecture, sovereignty. You also understand why I won't likely be able to carry this far. So I'm building in public and leaving breadcrumbs.

What Knowzy Is

At its core, Knowzy is simple: people ask questions, people answer them, sats flow to the respondents. But the architecture underneath is where it gets interesting.

Every answer you give becomes part of your dimensional profile. Over time, the system learns not just your demographics (age, location, income) but your preferences, values, and consistency. This isn't surveillance - you own this data. It's portable across nostr. You can reset it if you want, but you lose your history, your reputation, your accumulated value.

Question askers can target specific audiences: "Show this to suburban parents who previously said they care about privacy." The more specific the targeting, the higher the payment per response. Supply and demand, but with complete transparency about what dimensions are being used.

The consistency tracking is the key innovation. If you claim to be a fiscal conservative but consistently support increased government spending in polls, the system notices. Not to punish you - people's views evolve - but to create a marketplace for coherent perspectives. Over time, respondents with long, consistent histories become more valuable because their answers are predictive.

Why This Matters

Right now, your poll responses, survey answers, and stated preferences are scattered across dozens of platforms. Each one builds an index of you - but you don't control it, can't export it, and definitely don't benefit from it financially. Companies pay billions for market research while survey respondents get entered into drawings for gift cards. It's extractive.

Knowzy flips this. The dimensional profile you build is *yours*. The reputation you earn for thoughtful, consistent responses is *yours*. The sats you earn are *yours*. And because it's built on nostr, the whole thing is portable and censorship-resistant.

But more fundamentally: decentralized polling infrastructure matters because centralized polling has failed us. We've seen how platforms can suppress or amplify certain viewpoints, how survey framing shapes responses, how aggregators cherry-pick data to support predetermined narratives. When the indexes are transparent and the incentives are aligned, truth has a fighting chance.

The Architecture

I'm a data architect by training and temperament. I built a metadata-driven system of systems in ASP over 20 years ago - two pages, dozens of applications, still running in production (migrated to new tech only once). The principle was simple: get the data model right, make it flexible enough to evolve, and the application layer becomes almost trivial.

Knowzy follows the same philosophy. The schema I designed seven years ago still holds up. Questions, responses, user profiles, dimensional groups, tags, consistency tracking - it's all there in

the DDL files in the repo. I've added payment tables for Lightning integration and fields for nostr keys, but the core hasn't changed much.

The magic is in the DimGroup → DimType → MultiDim hierarchy. This lets you create arbitrarily complex audience segments without hardcoding anything. "Show me the answers from people who previously said X, Y, and Z" becomes a query, not a feature request. The metadata drives the targeting.

Starting with seed questions is crucial. New users answer 10-20 demographic and interest questions (and get paid for it) before they can access the broader question pool. This bootstraps their dimensional profile and ensures everyone in the system is immediately targetable. It also teaches the core loop: answer questions, earn sats, build reputation.

What I've Built (And What's Missing)

The database schema is solid. The conceptual architecture is sound. I've thought through the consistency algorithms, the payment flows, the anti-gaming mechanisms. The data model is in the repo at github.com/rjcraddock/knowzy.

What's missing is the implementation. The frontend doesn't exist yet. The Lightning integration is theoretical. The nostr event publishing is unwritten. The seed questions aren't authored. The consistency scoring algorithm is pseudocode in my head.

I've started this knowing I probably won't finish it. That's not defeatism - it's realism given my situation. But I'm sharing it anyway because incomplete work has value if it's documented. Someone else might see the architecture and understand what I was trying to build. Someone with more runway might pick it up and carry it forward.

For Whoever Comes Next

If you're reading this and thinking about building Knowzy (or something like it), here's what I'd suggest:

Start simple. The full vision is complex, but the MVP is straightforward: seed questions → user profiles → basic polling → Lightning payouts. Prove the core loop works before adding sophisticated targeting.

Trust the data model. I spent a lot of time on it. The flexibility is there - you can add dimensions, refine targeting, track new metrics without schema changes. Let the metadata do the work.

Consistency tracking is the moat. Anyone can build a poll app. The dimensional profiling over time is what makes this valuable. Protect the integrity of that data.

Build on nostr. Don't reinvent identity, key management, or social graphs. Use what exists. NIPs for zaps, relays for event storage, existing clients for reach. Integration over invention.

Keep the incentives aligned. The platform takes a small percentage to sustain itself, but the bulk of the value flows to respondents. If you optimize for extraction, you'll kill what makes this work.

The Index I'm Building

In my book, I argue that we've lost control of our own indexes - the categories, tags, and metadata that define how we're seen and sorted. Platforms index us for their benefit, not ours. Knowzy is an attempt to build a different kind of index: one that users control, that benefits them financially, that's transparent in its operation, and that prioritizes truth over engagement.

I won't see this finished. But I can leave the architecture, the schema, the thinking behind it. Someone else can build the index I won't live to use. That's enough.

The repo is public. The data model is there. The vision is documented. If you understand why this matters - why sovereign identity and aligned incentives and transparent indexes are worth building - then you have everything you need to start.

Build it better than I could have.

Jim Craddock
nostr.jimcraddock.com
www.redactedscience.org
npub: @Jim Craddock
github.com/rjcraddock/knowzy