

Luis Alfonso Reyes

AI Product & Delivery Engineer

Six systems. I built each one on my own, and five of the six are running right now — for a law practice, a manufacturer, a studio, and for paying subscribers.

luisreyesaldaco@gmail.com | autopoima.com/luis | tesseum.com | linkedin.com/in/luis-alfonso-reyes-aldaco-a30738122

Prague, Czech Republic | Czech permanent residency — EU work authorization, no sponsorship required

Each one says what it is, why it was built, how, and what changed — and carries its real status: **live, delivered** or **prototype**. I would rather you knew which than be impressed by the wrong one. Client names are withheld on purpose. Every figure was counted against the running system.

01 Tesseum

LIVE · PAYING SUBSCRIBERS

WHAT IT IS

A search and drafting tool for lawyers. You ask a legal question in plain language and it answers with the law behind the answer attached — the article, the ruling, the date — instead of a summary you then have to go and verify yourself.

WHY IT WAS BUILT

A lawyer working across more than one country has nowhere single to look the law up. It sits in a dozen government websites, in different languages, in documents built for printing rather than for reading by machine, and it changes every week.

A general AI assistant will answer the question and invent the citation. In law an invented citation is not a bad answer — it is malpractice, and it is the reason most firms will not let these tools near real work.

HOW IT WAS BUILT

I collected and cleaned the law myself: **736,833 articles from 6,537 laws across nine countries** and international bodies, plus **321,589 court decisions**. Everything is indexed by meaning rather than by keyword, so a question asked in Spanish will find the Czech ruling that answers it.

On top of that sits an assistant that drafts documents. Behind it sits the part that matters: before any answer leaves the system, every citation in it is re-checked against the documents that were actually used to write it. **A citation that cannot be confirmed is flagged, not shipped.**

I built all of it — the database, the back end, the interface, the billing, and the jobs that run every day to keep the law current.

IMPACT

A live business, not a portfolio piece: paying subscribers since October 2025, **1.47 million records** kept current, six tools in one product. The corpus is also sold as a data feed to other AI systems — 15 external clients, 1,048 requests served. One engineer, start to finish.

STACK React · Node · PostgreSQL with vector search · Anthropic, Google and Azure models under commercial contract · agent orchestration with LangGraph · Stripe · GDPR Art. 17 and 20 · automated evaluations gated in CI

HONEST SCOPE Docker and Compose, not Kubernetes. An AWS port of one component is packaged and not yet deployed, and I will say so until it runs.

02 Ventanas

DELIVERED · THE CLIENT RUNS IT

WHAT IT IS

The system a manufacturer uses to run a job from the customer's order to the finished window — and, for the first time, to know what that job actually cost.

WHY IT WAS BUILT

The company builds and restores historic windows. Its whole operation lived in thousands of loose spreadsheets and CAD drawings across shared folders. Nobody could answer "did we make money on that order?" without opening six files, and by then the material prices had moved. The data is commercially sensitive, so putting it in someone else's cloud was not an option.

HOW IT WAS BUILT

One application covering sales, production, purchasing, warehouse and shipping, with a catalogue that knows the materials and the hours each window needs. It reads their existing spreadsheets and drawings, so nothing had to be retyped.

It installs on their own server with a single command and upgrades itself without losing data — because I am not there to run it for them.

IMPACT

The number that mattered to them was cost, so I made cost behave. **An order's cost updates while the order is open and locks the moment materials are bought.** A price change next quarter can no longer rewrite last quarter's profit — which means that, for the first time, their margin history is worth arguing about. Running on their own hardware today.

STACK Python · PostgreSQL and DuckDB · Streamlit · Docker Compose · 14,365 lines, 87 automated tests, 16 versioned database upgrades

WHAT IT IS

The website and the paid online course of a movement studio — including the part that decides who is allowed to watch which video.

WHY IT WAS BUILT

The instructor had recorded a course and wanted to sell it herself. Every ready-made platform takes a large cut, and none of them stopped a paying customer from downloading the videos and passing them around. She also needed the whole thing in two languages, and needed to change her own class schedule without calling me every week.

HOW IT WAS BUILT

A 31-page site in Czech and English, plus a members' area. Customers sign in with a link sent to their email — no password to forget, lose or share — and each link works exactly once.

The videos are never handed over directly: every play request is signed and expires, so a copied address is dead on arrival. I tested that the way it should be tested, by trying to steal a video myself and confirming the server refused.

Payments run through the Czech bank gateway, and the customer's language survives the trip to the bank and back because it travels inside the part of the request the bank itself signs — so nobody can tamper with it on the way.

IMPACT

24 lessons behind a lock that holds. She runs her own schedule and her own campaigns from a spreadsheet and the site follows, which is the difference between a website and a dependency. The bank's compliance review came back with two conditions before go-live; I closed both, which meant writing her terms of sale, refund policy and complaints procedure. That part is not code, and it is why a lawyer who builds software is useful.

STACK Serverless Node on Vercel · PostgreSQL · signed video delivery · passwordless single-use tokens · Czech bank gateway with RSA-signed requests · Czech and English throughout

04 Marca

WHAT IT IS

The public site of a trademark lawyer — the first thing a prospective client sees, and the place where they can start a case without having to phone anyone.

WHY IT WAS BUILT

She had a printed brochure and no presence online at all, in a practice where clients start by searching. She also had no way to answer, at two in the morning, the only question every visitor actually arrives with: what does this cost.

HOW IT WAS BUILT

I built the site out of her own brochure — its colours, its typefaces, its prices — so a client who has both in front of them sees one firm and not two. It carries an assistant that answers the common questions about registering a trademark, and an intake split into two short steps instead of one long form.

It ships in two complete looks, light and dark, chosen by the visitor: every colour is defined by its job ("this is a background", "this is a heading") rather than by its value, so switching is a swap and not a second website to maintain.

IMPACT

The first draft repeated her six practice areas three times over. **I cut the page by roughly a third — 1,027 words down to 809 — while adding the price list that had never existed.** Shorter, and it finally answers the question people came with.

STACK Plain HTML, CSS and JavaScript, no framework, on purpose · role-based colour tokens · serverless assistant endpoint · contrast verified against rendered pixels, not against the code

HONEST SCOPE The intake form validates but does not yet submit anywhere. That is the client's next step, not a thing I left broken.

WHAT IT IS

An experiment on my own legal corpus, testing whether a search engine can be made to find the case it should have found.

WHY IT WAS BUILT

Search by meaning has a blind spot that nobody advertises. It finds documents that sound like your question. **It cannot find the landmark ruling that every relevant case cites but that shares almost no words with what you typed** — and in law that is very often the one you needed. I wanted to know whether the citations themselves could close that gap.

HOW IT WAS BUILT

I rebuilt part of my corpus as a network of who cites whom: 622 court decisions, 4,063 citations between them, and 4,292 pieces of legislation with 1,393 internal cross-references. A search now runs twice — first by meaning, then by walking outward along the citations from whatever the first pass found.

IMPACT

On a test question about enforced disappearance, ordinary search returned five cases. Walking the citations surfaced **two landmark judgments it had missed completely**. The same exercise exposed 24 corrupted citations that had been sitting unnoticed in the source data — a data-quality problem I would never have seen otherwise.

STACK Neo4j · Cypher · the same embedding model as the production corpus · runs end to end in one to three seconds

HONEST SCOPE Days old, a command-line prototype, not yet inside the product. Built and running on real data — not years of experience with graph databases.

06 Slovo

LIVE SINCE 30 JULY 2026

WHAT IT IS

The system that answers a movement studio's Instagram followers on her behalf. She posts "*comment the word DECH and I'll send it to you*"; when somebody does, the material arrives in their private messages within seconds, and their comment gets a public reply. Same studio as 03. *Slovo* is Czech for "word" — it is the word she asks people to write.

WHY IT WAS BUILT

The offer only works if it is instant, and she was doing it by hand: reading every comment, opening each person's messages, pasting the link. On a post that lands, that is a hundred conversations in one evening, and the people she got to the next morning had already moved on.

She teaches most of the day, so the hours when a post is hottest are exactly the hours she cannot answer. And she did not want to call a developer every time she changed a word.

HOW IT WAS BUILT

Everything she controls lives in her own spreadsheet: the trigger words, the private message, the public replies — the bot rotates between them at random, so the comment section does not read like a machine — and an on/off switch per campaign. It re-reads that sheet every five minutes. Nothing to restart, nobody to call.

The harder half was deciding which posts to watch. Instead of registering each campaign by hand, the bot reads her captions and recognises her own call to action: the word *slovo* followed closely by a trigger word, which is how she has always written it. Tested against 47 of her real posts, it found both campaign posts and left alone the other 34 that merely mention the same words — including the trap, a post that says "write to me in the comments" and also says the word, but never asks anyone to comment it.

It is built to fail closed. With no trigger words configured it does nothing. If it cannot tell which posts are campaigns it acts on none — it never falls back to answering everything. Words match at word boundaries, so it catches the Czech inflections (*dechu, dechem, během*) without firing on *průběh*, where somebody was simply asking a question. It never replies to itself, and it never goes back to message people who commented before the campaign existed.

The message text is never written to the log. Her clients write about their bodies, their health and their appointments; that is sensitive data under GDPR and it has no business sitting in a server log. Comment text is kept, because it is public and it is what you need to debug a trigger word.

IMPACT

Live since 30 July 2026 and still running. In its first six weeks it answered **161 different people — 173 private messages and 174 public replies**, every one of them within seconds, none of them typed by her. A new campaign is picked up on its own within fifteen minutes of her posting, so she can launch one without telling me. That was the point: not to automate a task, but to remove myself from it.

STACK Python · Instagram Graph API webhooks · Google Sheets as the control surface · systemd on my own always-on server · public endpoint over Tailscale Funnel · automatic token refresh

HONEST SCOPE One failure mode I have not fully removed: if the platform token expires the bot stays alive and the sends fail quietly. There is an automatic refresh and a status file, and I still check it.

THE CAREER THAT PAID FOR IT

Oct 2025 — present

Founder & AI Engineer

Tesseum · Prague

Jun – Oct 2025	Forecast Analyst	Groupe SEB · Prague — annual demand forecast built end to end in Python, SQL and Excel
May 2021 – May 2025	Volume Management Analyst	ExxonMobil · Prague — SAP volume management for a global operation; process automation; S&OP executive forums
Sep 2020 – Apr 2021	Order Specialist	TP Vision · Prague — order-to-cash in SAP SD, external customers, Salesforce
2019 – Aug 2020	Manager	A five-restaurant group · Prague — ~12 staff per shift; owned costs, pricing and margin
2011 – 2017	Legal Counsel	Private practice · Mexico — civil, commercial and corporate; counsel to senior management

EDUCATION, LANGUAGES & PUBLICATION

- **Master of Law (LL.M.)** — Charles University, Prague, 2019
 - **Licenciatura en Derecho** — Mexico
 - **Certified Ontological Coach** — Newfield Network
- **Languages:** Spanish (native) · English (C2) · Czech (B2) · Portuguese (B1)
 - **Author** — *Epigénesis biológica del Derecho / The Human Code*, Amazon, January 2026. Ten years of independent research in legal philosophy.