

KATYELLA

AI-Powered Document Automation for Specialty Pharma

How Katyella Delivered a Production-Ready AI Platform
in 4 Weeks — With Almost Zero Meetings

Client: Evo Advertising (via Lucien Vaccaro)

Industry: Specialty Pharma / Healthcare Market Access

Engagement: Fixed-Scope MVP Development & AI Integration

Duration: 4 Weeks

KEY RESULTS AT A GLANCE**75%**Time Reduction
(4 hrs → 30 min per
letter)**100%**Clinical Claim
Coverage**4 wk**SOW to Production
Delivery**<\$500**Monthly Infra
Cost

The Challenge

Evo Advertising's pharma client faced a critical bottleneck in their market access workflow: **Letters of Medical Necessity (LMNs)**.

Every time a physician needed to prescribe a specialty gene therapy, they had to manually draft a multi-page letter to the patient's insurance payer — citing specific clinical evidence, aligning to payer-specific coverage criteria, and formatting everything to regulatory standards.

The problems were compounding:

- **Hours per letter.** Each LMN required physicians to research payer policies, pull clinical citations, and draft custom prose. A single letter could take 2-4 hours.
- **Payer fragmentation.** Every payer (Aetna, UnitedHealth, Cigna, BCBS, etc.) has different coverage criteria, required documentation, and policy language. One template doesn't fit all.
- **Compliance risk.** Manually assembled letters risked omitting required clinical claims or using non-approved language.
- **Provider friction.** The administrative burden discouraged providers from prescribing — directly impacting patient access to treatment.

Evo needed a working demo platform to prove that AI could solve this problem — not a slide deck, not a wireframe. A real, functional tool their pharma client could test.

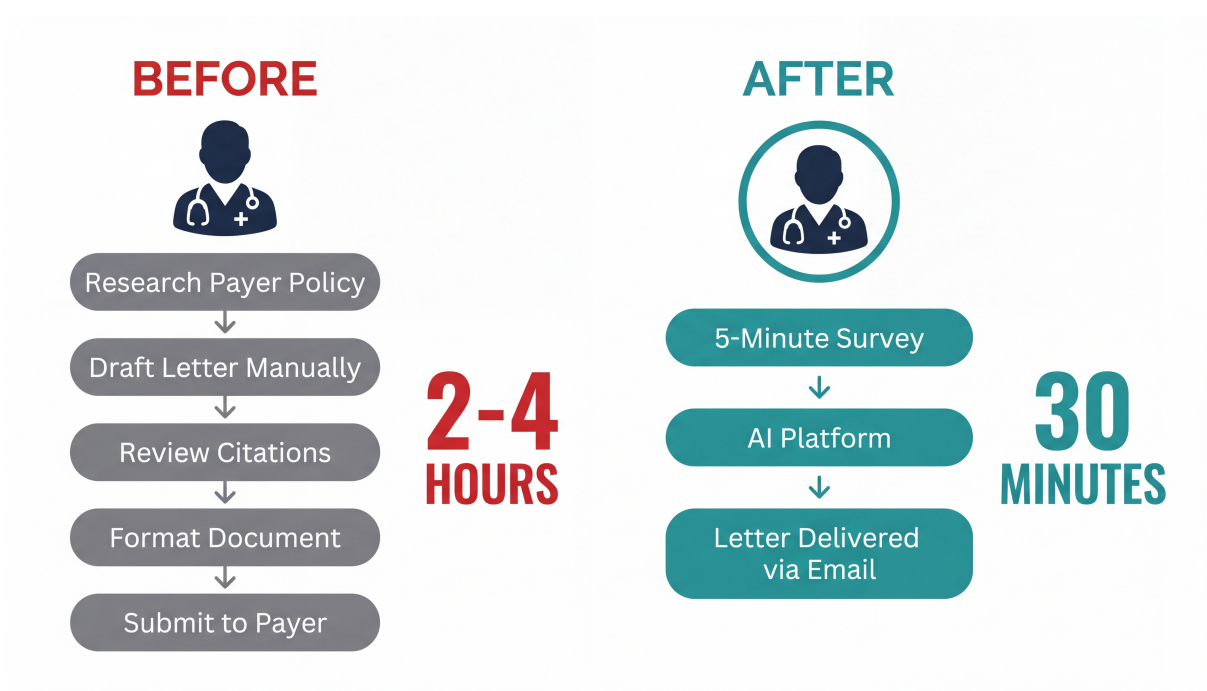


Figure 1: Before vs. After - Manual LMN Drafting vs. AI-Powered Automation

The Solution

Katyella built a **full-stack AI platform** that transforms a 5-minute physician survey into a complete, payer-compliant Letter of Medical Necessity — generated in under 30 minutes.

How It Works

- 1. Guided Survey (29 clinical data points)** A structured form walks the physician through 6 sections: coverage details, clinical profile, functional assessments, treatment history, quality-of-life impact, and clinical rationale. Auto-save ensures no data loss.
- 2. Deterministic Template Assembly** Survey answers are routed through a logic map that selects the correct pre-approved clinical claims for the specific payer and authorization type. This is not prompt engineering guesswork — every claim is deterministically selected from a curated library.
- 3. AI Polish Pass** Claude processes the assembled draft, improving grammar, flow, and professional tone — without removing or altering any clinical claims. If the AI pass fails for any reason, the system falls back to the deterministic draft. Zero data loss.
- 4. Dual-Format Document Generation** The platform generates both an editable Word document (for physician review and customization) and an archival PDF — simultaneously. PHI placeholders appear in highlighted brackets for easy identification.
- 5. Email Delivery** The completed letter is delivered to the physician's inbox within 60 seconds of submission, with a direct download link.
- 6. Admin Backend** A searchable, filterable dashboard tracks every submission with a complete audit trail: survey responses, generated letter text, document URLs, claims used, and timestamps.

AI Document automation platform

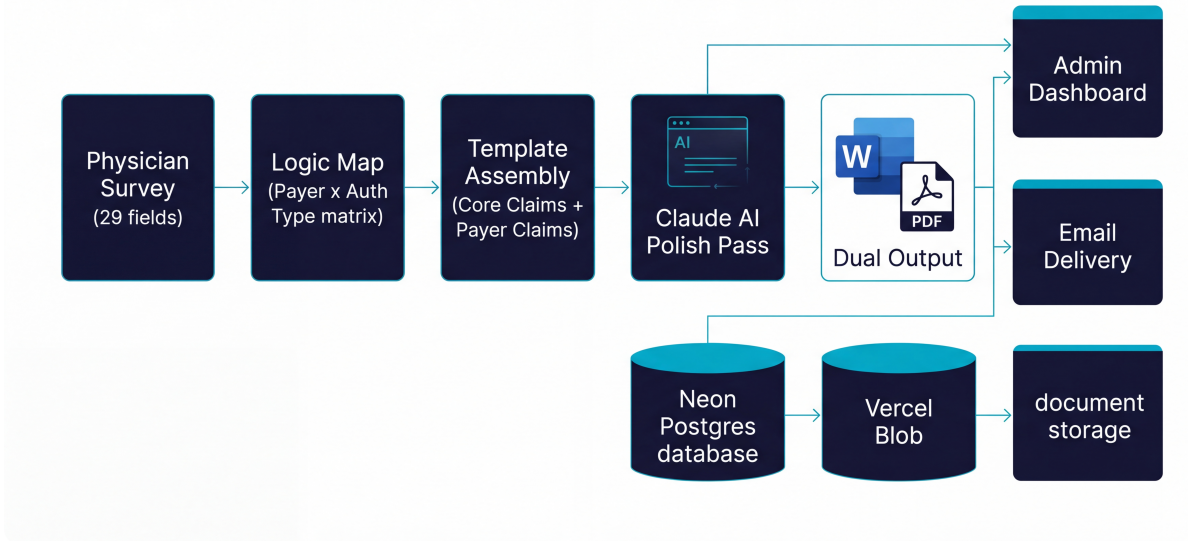


Figure 2: System Architecture - End-to-End Document Generation Pipeline

Payer Intelligence

The platform supports a matrix of payer-specific claim libraries:

	Initial Authorization	Reauthorization
Aetna	8 payer-specific claims	7 payer-specific claims
UnitedHealth	9 payer-specific claims	7 payer-specific claims

Plus 29 core clinical claims shared across all payers – each derived from approved clinical evidence with citations.

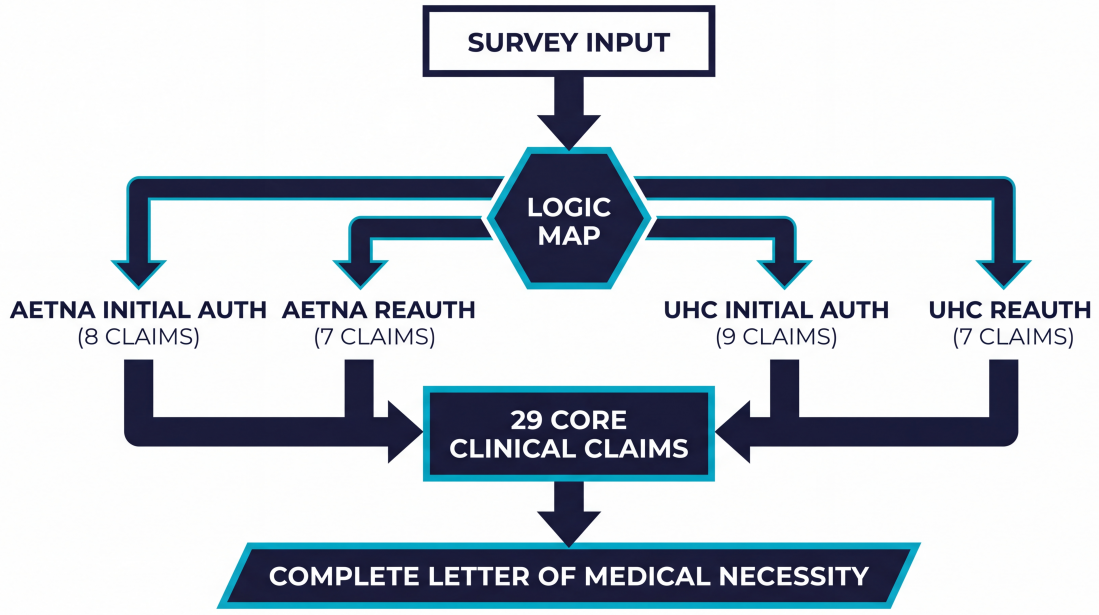


Figure 3: Payer Intelligence Matrix - Coverage Criteria by Payer and Authorization Type

Tech Stack

Layer	Technology
Frontend	Next.js 16, React 19, TypeScript, Tailwind CSS 4
AI Engine	Claude API (Haiku 4.5) via Anthropic SDK
Database	Neon Serverless PostgreSQL + Drizzle ORM
Documents	Custom .docx and .pdf generation (docx + pdf-lib)
Email	Resend (transactional)
Storage	Vercel Blob
Hosting	Vercel (auto-scaling, edge deployment)
Testing	Vitest (unit) + Playwright (E2E)

Monthly infrastructure cost: Under \$500/month at production scale.

The Delivery

4 Weeks. Minimal Meetings.

Week	Milestone
1	Requirements finalized, clinical content ingested, survey form built
2-3	AI integration, letter generation pipeline, document output, email delivery
4	Admin backend, end-to-end testing, staging deployment, client review

What the client expected: Weeks of discovery meetings, multiple design revisions, iterative feedback loops to nail the look and feel.

What actually happened: Katyella shipped full working prototypes early and often, paired with async video walkthroughs. The client could test real functionality — not review mockups. Style and UX questions resolved themselves because the client could *use the thing* instead of imagining it.

The result: a production-ready platform delivered on schedule with almost no synchronous meetings and minimal revision cycles.

Professional 4-week Project Delivery Timeline

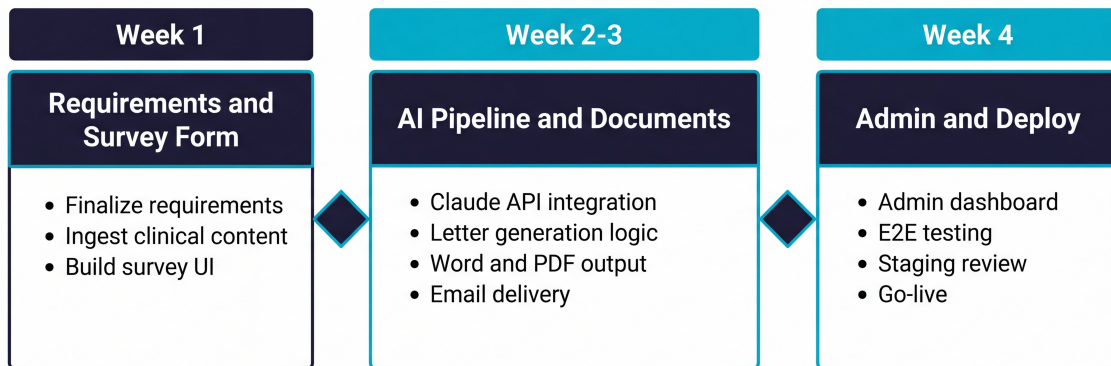


Figure 4: 4-Week Delivery Timeline - From SOW to Production-Ready Platform

Results

Speed

- **Before:** 2-4 hours per letter (manual physician drafting)
- **After:** Under 30 minutes per letter (AI-generated from survey)

Quality

- 100% of required clinical claims included in every letter (deterministic routing, not probabilistic)
- Payer-specific policy alignment baked into the logic — not left to the physician's memory
- Complete audit trail for every submission

Delivery Efficiency

- **4-week delivery** from signed SOW to working platform
- **Minimal back-and-forth** — working prototypes replaced meetings
- **Async-first communication** — video walkthroughs instead of status calls

Business Impact

- Demo platform led to an expanded engagement with the client
- Proved the viability of AI-assisted document automation in a regulated industry
- Reusable architecture supports additional drug products, payers, and authorization types

Why It Worked

1. Prototypes Over Presentations

Instead of burning weeks on wireframes and design reviews, Katyella shipped functional prototypes within days. The client tested real software, not slide decks. This collapsed the feedback loop from weeks to hours.

2. Deterministic AI, Not Prompt Roulette

The letter generation pipeline doesn't rely on AI to "figure out" what to include. A logic map deterministically selects every required claim. AI only handles prose quality — grammar, transitions, tone. If the AI layer fails, the letter still generates correctly from the template. This isn't a chatbot — it's an engineered system with AI polish.

3. Full-Stack, One Team

No handoffs between a design agency, a frontend shop, and a backend contractor. Katyella owned the entire stack — from the survey UX to the AI pipeline to the document generation to the email delivery to the admin dashboard. One team, one codebase, one deployment.

4. Right-Sized Architecture

Serverless Postgres, edge-deployed Next.js, blob storage, transactional email. No over-engineered microservices. No Kubernetes. Infrastructure costs under \$500/month — not \$5,000. The architecture matches the actual scale of the problem.

About Katyella

Katyella LLC is a software consultancy specializing in AI-powered application development, Java modernization, and data pipeline automation. We build production systems — not proofs of concept that need to be rebuilt.

Ready to automate a manual workflow? [Contact us](#)