

Konrad Odell

Staff Engineer, Architecture · Applied AI & Healthcare Data

Dallas-Fort Worth, TX · contact@konradodell.com · konradodell.com · github.com/konradish · linkedin.com/in/konrad-odell

SUMMARY

I've been shipping software for twenty years and setting technical direction since 2018. I'm on the architecture team at Moxe Health, a HIPAA-regulated clinical data exchange (joined 2022, architecture team since 2024), working on the Databricks lakehouse, FHIR pipelines, and evaluating LLM and agent approaches against real clinical and billing workloads. On my own time I built and run ElectricityFinder, which runs an LLM extraction pipeline in production in a regulated market, and the tooling I build with: MCP servers, agent skills, and evaluation harnesses. I build with coding agents in a worker and verifier loop: a separate agent with fresh context has to prove the work runs before I accept it.

EXPERIENCE

Staff Engineer, Architecture, Moxe Health

Jan 2024 – Present

Architect for a B2B healthcare data interchange platform (HIPAA-regulated clinical data at scale), focused on data-platform capability, applied-AI evaluation, and cross-team technical coordination.

- Identified, scoped, and validated applied-AI opportunities against real clinical and financial workloads, including a multimodal LLM proof of concept that recovered \$300K in invoices that would otherwise have taken manual review to find
- Published a Claude Code skill that diagnoses production incidents from Grafana logs and traces, including a Jira template that requests the production observability access it needs; adopted by multiple teammates
- At an executive's request, evaluated moving a production document-ingestion LLM pipeline to another model provider: benchmarked alternatives on sampled documents and documented the migration path
- Designed a FHIR-to-template clinical document generation pipeline
- Partnered on the Databricks lakehouse (Unity Catalog governance, self-service analysis tooling) and coordinated architecture reviews for a large-scale clinical lakehouse migration off a legacy external provider
- Coordinated architecture decisions across data platform, clinical retrieval, and revenue-operations functions, presenting tradeoffs to engineers and senior leadership; coached senior engineers on architectural thinking and mentored across disciplines
- Driving canary deploys, smoke tests, and auto-rollback as first-class patterns in the new platform architecture (in progress)

Senior Software Engineer / Team Lead, Moxe Health

Sep 2022 – Jan 2024

Led the consent-management service team on a HIPAA-regulated B2B clinical data interchange platform; owned delivery and deploy safety, and co-owned the service architecture.

- Led consent-management service team (Java, Spring Boot, NATS) with zero-downtime blue/green deployment
- Ran the team's delivery cadence and code reviews

Senior Software Engineer II, Cognizant Softvision → Discover Financial Services

Aug 2021 – Sep 2022

Built card-tracking API ingesting USPS and internal events on Spring Boot microservices in a regulated financial-services environment.

- Introduced contract-test pipeline (WireMock + Pact) to catch breaking API changes pre-production
- Partnered with DevOps on Helm-based release process, reducing deploy time from 40 to 8 minutes

Senior Software Engineer → VP of Research, Nearstar Inc.

Sep 2010 – Aug 2021

Legacy modernization, Docker adoption, CI/CD implementation, and performance optimization; grew from senior IC to leading technical direction as VP of Research (2018-2021).

- Dockerized monolith suite enabling per-customer configuration and a SaaS pivot
- Optimized multithreaded PDF loader during COVID spike: +260% throughput, backlog cleared in 36 hours

Software Developer, ClubCorp

Jun 2001 – Jul 2008

Mainframe to Oracle migration; built authentication framework with salted password hashing.

PROJECTS

ElectricityFinder · electricityfinder.online

Solo-built service that helps Texans find an electricity plan: identified the niche, wrote the spec, obtained the broker license, and shipped it to production in a regulated domain.

- LLM plan-document extraction pipeline with checks at ingest: extractions that fail cross-contamination or confidence checks are rejected or flagged before they can be priced
- Full stack solo: Cloudflare Workers + D1, React/TypeScript frontend, Smart Meter Texas usage integration
- Registered PUCT electricity broker (BR260062); PUCT disclosure, CAN-SPAM, and TDPSA retention compliance implemented and tested
- ~5,000 automated tests, deploy-time bundle invariants, and post-deploy content smoke checks that assert rendered results, not just HTTP status

Agentic Engineering Estate

A working multi-agent development system used daily to build and verify the projects on this page — the practice behind treating AI coding agents as production build tools.

- Reusable building blocks: a library of 68 agent skills (usage-instrumented and pruned by actual fire-count), 45 deterministic hooks, MCP servers and tool adapters in daily production use, and sub-agent/workflow orchestration patterns
- Custom eval harnesses rather than vibes: an own-commit SWE-bench-style benchmark for scoring local and frontier models on real repository work, plus adversarial-eval and domain-specific benches. Caught a model reading the answer from git history, closed the leak, and re-audited all 122 prior runs
- Deterministic orchestration with worker/verifier separation: unattended overnight agent runs fixed 8 ElectricityFinder bugs, each re-verified from a clean checkout by a separate agent; deploys stayed human-approved
- Belief-state engine pairing LLM extraction with deterministic code judges, over a hybrid lexical + semantic search layer
- Multi-provider routing in practice (Anthropic, AWS Bedrock, Google, and local open-weight models behind one gateway), plus local serving and LoRA fine-tuning on consumer hardware

SKILLS

Applied AI & Evaluation: Production LLM pipelines, AI coding agents as production build tools, Multi-agent orchestration, MCP servers & tool integrations, Reusable agent skills / sub-agents, Retrieval & hybrid semantic search, Multi-provider routing (Anthropic, AWS Bedrock, Google, local open-weight), Custom eval harnesses, LLM-as-judge gates, Worker/verifier separation, Local inference & LoRA fine-tuning

Healthcare & Regulated Domains: HIPAA, FHIR, quality measures, Clinical data processing, PII governance, PUCT / CAN-SPAM / TDPSA

Languages: Python, TypeScript, Java, SQL, JavaScript, Bash

Data, Cloud & Delivery: Databricks (Unity Catalog, Genie, Lakehouse), Iceberg, DuckDB, Vector databases, AWS (API Gateway, Lambda, EKS, RDS, S3), Cloudflare (Workers, D1, Tunnels), Event-driven architecture (NATS), Docker, Kubernetes, GitHub Actions, Blue/green deploys

Working familiarity: Trino, dbt, HL7 v2, Supabase

EDUCATION

B.S. Computer Science, UT Dallas, Cum Laude (Aug 2010)

WRITING

I Built a Factory Control Plane for AI Agents in One Afternoon · konradodell.com/blog/dispatch-factory-dashboard.html (Mar 2026)

Stop Getting Attached to Your AI Tooling · konradodell.com/blog/compostable-patterns.html (Feb 2026)

The Rutter Pattern: AI Memory Through Continuous Refinement · konradodell.com/blog/rutter-pattern.html (Dec 2025)