

DAVID VEKSLER

Head of AI / Director of AI Engineering · Frontier-LLM (Claude) deployment in regulated organizations

Denver, CO · +1 214 659 1775 · davidleovexsler@gmail.com

linkedin.com/in/davidveksler · github.com/davidveksler · davidveksler.com/work/regulated-lender-ai-platform/

SUMMARY

I founded the AI delivery function at a regulated commercial lender, and I will do it for you without breaking your system of record. I lead with platform and governance before breadth: a skill marketplace as the distribution backbone, a publish-and-audit gate, human-in-the-loop as an architectural property, observability with cost-per-skill, and fair-lending (ECOA/FCRA) and model-risk (SR 11-7) framing for anything credit-adjacent. Twice, in two settings: a founding build at a regulated lender, and enterprise Claude adoption at a Mars diagnostics company. I move executive conversations from opinion to arithmetic. US citizen; remote, or hybrid in Denver.

EXPERIENCE

SPECTRA · Principal AI Engineer, founded the AI delivery function (*embedded contract, concurrent with Antech*)

Regulated CRE bridge-lending fund, ~3 dozen people, 9 departments · Spring 2026 · reported to CTO, CEO sponsor · case study: davidveksler.com/work/regulated-lender-ai-platform/

- Stood up the entire AI delivery function from zero across nine departments: a CI-distributed **10-plugin skill marketplace** (~30 skills, 19 firm-authored, 6 populated namespaces, 8 connectors), a **publish-and-audit governance pipeline**, a chat-to-ticket helpdesk, a scheduled enrichment service, and AI observability.
- **Governance as load-bearing, not paperwork:** a skill auditor runs deterministic checks (secrets and PII scan, schema, connector-declared-equals-used) plus an LLM judgment pass that routes fair-lending, financial-calculation, system-of-record-write, and NPI flags to named reviewers, and hard-rejects external-send-without-HITL. It produces a report, not a verdict; a human merges.
- **Observability from zero:** OpenTelemetry into a wide-event backend, cost-per-skill via a session-id join, and a who-did-what-when history that partially closed a platform compliance blind spot.
- **The architecture spine that kept the AI layer out of the CTO's blast radius:** the system of record always wins (AI reads and drafts, never overwrites); drafts-only by default, autonomy a per-skill dial earned from telemetry rather than asserted; the connector credential, not the marketplace, is the real boundary. Threat-modeled confused-deputy and prompt-injection risk, with the human checkpoint one layer below the agent, never inside it.
- **Measured impact:** a multi-hour manual pre-quote sweep collapsed to a single multi-minute run; adoption instrumented per skill, not projected. Ambiguous mandate to signed-off MVP in 16 days. Clean exit on a fund redemption.

ANTECH DIAGNOSTICS (Mars / VCA) · Principal Engineer & internal Applied AI Architect

Remote · Nov 2024 – Present

- De facto internal AI architect for a regulated diagnostics org of 50+ engineers: drove Claude adoption from evaluation through governed deployment and ran the enablement that produced unprompted adoption across teams. Proof the founding pattern is repeatable, current, and at enterprise scale.
- Built the business case that moved leadership from skepticism to a sponsored push for an organization-wide Claude agreement, and shipped the governed skill catalog underneath it (reviewed, versioned, read-only by default, human approves production, review/eval gate, per-skill observability) so the org adopts on one auditable foundation instead of scattered private use.
- Led the **7-engineer team** that built a support observability platform end to end with agentic workflows on Claude Code, and put LLM-driven test generation behind a CI gate: coverage 5% to 80%+ in one sprint, and **sprint throughput 4.2x** against the same team, sizing, and backlog.
- Decomposing a 20-year SOAP/Java monolith into .NET 8 microservices on 50K+ daily diagnostic orders; MTTR weeks to under 24 hours, uptime 97% to 99.99% in the same quarter.

LEADERSHIP & REGULATED-DOMAIN TRACK RECORD

- **Celsius** · Senior Software Engineering Manager (2022–2024): directed three distributed teams, **24 engineers** across the US and Europe, under active regulatory scrutiny: hiring, allocation, performance calibration, and architecture and risk communication to executives. \$100M+ TVL; Chainlink proof-of-reserve; \$4M+ recovered for creditors during the wind-down.
- **Royalty Exchange** · CTO (2020–2021): AWS migration, PCI-DSS readiness, world-first music-royalty NFT platform (Billboard); weekly releases; mentored three engineers into lead roles.
- **Vellum Capital** · Co-Founder (2018–present): CCSS-L3 custody (MPC, sub-1hr DR); hedge-fund infrastructure.
- **FEE** · Director of Technology (2015–2020): \$3M platform, full SDLC, recruited the team.
- Earlier: **Liberty.me** (Founder & Principal Architect) · **EF Education First** (Software Architect) · **Match.com** (Senior Application Engineer).

PUBLIC PROOF

Governing Agentic AI, a field guide on the substrate-not-agent thesis and four invariants, built by the governed pipeline it describes: 160+ pages from one checked-in spec, human merge gate, 770+ public commits. cheatsheets.davidveksler.com/governing-agentic-ai.html

SKILLS

AI platform strategy and delivery, enterprise AI adoption and advisory, frontier-model evaluation to production (Claude, multi-vendor), LLM eval-framework design, AI governance and human-in-the-loop (ECOA/FCRA, SR 11-7, NPI/PII), agent and skill architectures and marketplaces, AI observability and cost-per-skill, legacy modernization, executive stakeholder management, engineering management (24 engineers), C#/ .NET, Python, TypeScript/React, Azure, AWS, SQL, OpenTelemetry, CI/CD

EDUCATION

M.S. Management Information Systems · B.S. Economics & Political Science, Texas A&M · Data Science Immersive, General Assembly