Rob Dalida

AI Engineer III | Production AI Systems | Agentic Workflows | AI/ML Engineering

Email: robertdalida@gmail.com | LinkedIn: linkedin.com/in/robertdalida | GitHub: github.com/robertdalida | East Coast, UTC-5

AI Engineer with 8 years in operations and 2 years in AI/ML delivery, building production AI systems end to end across data, model behavior, deployment, monitoring, and cross-functional adoption. Background in school-based Medicaid operations gives a practical lens for designing AI systems that remove manual workflow burden, improve compliance, and turn ambiguous business needs into reliable engineering patterns.

CORE SKILLS

AI/ML Systems: LLM applications, RAG, agentic workflows, evaluation harnesses, prompt/context engineering, semantic embeddings, model behavior monitoring, AI observability.

Engineering: Python, TypeScript, FastAPI, Next.js, React, ReactFlow, D3.js, PostgreSQL, Snowflake, GitHub Actions, CI/CD.

AWS: Lambda, Bedrock, SageMaker, S3, EventBridge, SES, CloudWatch, Cognito.

Operating Strengths: Architecture tradeoffs, production ownership, cross-functional delivery, security remediation, compliance workflows, stakeholder translation.

SELECTED IMPACT

- Reduced initial-prompt token usage by 59.3% by redesigning a broad agent into an orchestrator and specialist-agent pattern while preserving a single assistant experience.
- Improved AI answer coverage from approximately 23% to 80% through a governed auto-research loop that proposes bounded changes, reruns deterministic evaluation, and keeps only improvements.
- Reduced security findings from 468 to 170 in 5 days by integrating Snyk MCP with Codex-based triage and repo harness workflows; cleared all 10 critical findings.
- Improved RAG relevancy from 0.912 to 0.976 across 30 support scenarios using Resolution Cards built from solved operational knowledge.
- Shipped 11 production systems across AI infrastructure, compliance automation, workflow tooling, knowledge systems, and cross-functional operational support.

EXPERIENCE

AI Engineer III - Frontline Education

Customer Experience AI | 2025-Present

- Own production AI systems from source data through model behavior, deployment, monitoring, and iteration after launch.
- Drive architecture decisions for reusable AI patterns, evaluation workflows, repo harnesses, and agentic execution paths.
- Partner across Support, Product, Legal, GTM, Services, and Security to translate business pressure into practical AI systems.
- Standardize practices for source-grounded knowledge systems, retrieval evaluation, agent orchestration, and operational AI review.

- Build systems that reduce manual workflows, improve auditability, and make AI outputs easier for teams to inspect and trust.

Medicaid Services Manager - Frontline Education

School-based claiming and operations | 2022-2025

- Owned IEP documentation, Medicaid claiming workflows, compliance support, and client operations.
- Led service operations across complex Medicaid workflows and helped translate client needs into repeatable operational patterns.
- Built the domain foundation that later shaped AI engineering work focused on practical workflow removal and operational reliability.

Medicaid Client Manager - Frontline Education

Medicaid services and client operations | 2018-2022

- Managed Medicaid client relationships and day-to-day operational support for school-based claiming programs.
- Developed expertise in compliance-heavy operational workflows, service delivery, and software adoption pain points.

SELECTED PROJECTS AND SYSTEMS

Content Graph - AI Knowledge Graph for Support Content

Built an AI knowledge graph pattern that embeds learning content, identifies missing links, and surfaces coverage gaps to improve retrieval and support deflection.

Flow Builder - Visual Data Transformation Workflow Tool

Built a ReactFlow-based workflow builder for no-code-style data transformation and calculation pipelines.

Model Workbench - Machine Learning Platform

Built machine-learning platform concepts for expense classification, including datasets, training/evaluation flow, model registry patterns, and classification testing.

Rules Engine - Compliance Automation Platform

Built a nationwide compliance-engine pattern that sources regulatory documents, translates requirements into structured outputs, and helps product and engineering create business rules.

Signal Console - AI Observability Surface

Built observability concepts for AI support systems, including anomalies, retrieval behavior, guardrails, incidents, and team-facing operational review.

Pipeline Scoring - GTM and Sales Intelligence

Built pipeline intelligence concepts for ranking accounts, identifying expansion opportunities, and highlighting clients with upsell potential.

PII Redactor - Production Compliance Automation

Built production Zendesk SSN redaction workflows with scanner/redactor phases, dry-run safety, backups, daily Lambda automation, and stakeholder reporting.

COI Automation - Legal Workflow Automation

Built an AWS-based Certificate of Insurance pipeline with secure storage, scheduled processing, audit trails, personalized email delivery, and PDF generation.

Repo Agentic Harness - AI-Assisted Engineering Contract

Created a reusable repository contract for AI-assisted engineering, including entrypoints, ownership boundaries, verification commands, secret stop conditions, and handoff rules.

WRITING AND TECHNICAL NOTES

Self-Maintaining Knowledge System: Source-grounded markdown wiki pattern for keeping internal documentation structured, reviewable, and easier to improve over time.

Codex Skills as SaaS-Style AI Product: Repo-wiki workflow that turns real repositories into browseable, refreshable documentation without a separate SaaS tool.

AI Auto-Research Loop: Governed loop for proposing one bounded knowledge-base change, scoring it, and keeping only deterministic improvements.

Resolution Cards: Operational knowledge pattern for improving AI support response relevancy and faithfulness.

Harness Engineering: Repository design pattern for AI coding agents with clear entrypoints, workflow contracts, verification paths, and stop conditions.

Context Engineering: Retrieval pattern for giving LLMs the right slices of context instead of overwhelming them with long, noisy prompts.

RECOGNITION

- Recognized by senior LMS, App Security, Strategic Initiatives, GTM Strategy, and Legal partners for practical workflow design, security remediation, automation delivery, and cross-functional follow-through.