

Keith Steward, Ph.D.

Principal Applied AI Engineer | Production Agents, Evaluation & Platforms

Boston, MA | 508-318-8035 | keith.steward@gmail.com

[LinkedIn](#) | [IntelligentDataWorks.ai](#) | [IDW LinkedIn](#)

PROFILE

Hands-on applied AI engineer and architect who turns ambiguous model and agent ideas into reliable production systems. Built the core and first 80% of an agentic automation platform personally, then scaled it to 70+ automations across more than 20 workflows with 99.9% availability. Combines Python, APIs, data workflows, RAG, evaluation, observability, security, cloud infrastructure, and pragmatic technical judgment with a record of translating research and customer needs into reusable software.

SELECTED CAREER HIGHLIGHTS

- Built AssistX hands-on with Python, FastAPI, agents, RAG, PostgreSQL/PGVector, APIs, Docker, AWS, and CDK; implemented the core and first 80% of its initial application before expanding the engineering team.
- Shipped and operated 70+ production automations across more than 20 workflows, accelerating selected work by up to 50x while sustaining 99.9% availability.
- Established evaluation, LLM-as-judge, guardrails, failure analysis, observability, auditability, model routing, cost controls, CI/CD, release, and incident-response practices for production AI.
- Advised more than 200 enterprises on AI/ML, data platforms, architecture, cloud adoption, and production delivery; created 170+ technical enablement assets and reusable implementation practices.
- Delivered cloud data, privacy, analytics, and machine-learning systems in regulated enterprises, including a privacy platform that automated approximately 90% of a manual workflow.
- Works effectively across research code, maintainable services, deployment automation, and customer-facing product experiments; uses instrumentation and production evidence to decide when to iterate quickly, invest in reusable infrastructure, or choose a simpler implementation.

PROFESSIONAL EXPERIENCE

Technical Co-Founder & CEO | Intelligent DataWorks | Boston, MA | Aug 2023 to Present

Lead an AI engineering organization while remaining hands-on in platform architecture, implementation, delivery, governance, and operations.

- Recruited and led 11 engineers selected from more than 400 applicants; established onboarding, architecture training, Agile sprints, mentoring, code review, delivery priorities, and production standards.
- Built the AssistX core and first 80% of its initial application hands-on using Python, FastAPI, LLMs, agents, RAG, PostgreSQL/PGVector, APIs, Docker, AWS, and CDK.
- Delivered 70+ automations across more than 20 workflows, accelerating selected work by up to 50x while sustaining 99.9% availability.
- Established AI evaluation, LLM-as-judge, guardrails, observability, auditability, model routing, cost controls, security, CI/CD, release, and incident-response practices.
- Own a \$900,000 budget, staffing, product and technical roadmaps, architecture, releases, cloud operations, and executive decisions across the venture.
- Initiated strategic acquisition discussions regarding IDW's AI platform as part of a planned transition to a senior AI/engineering leadership role.

Sr. Director of Engineering | SuccessKPI | Remote | Jun 2022 to Aug 2023

Built and led engineering for a new AI-powered workforce-management platform within a cloud-native enterprise SaaS company.

- Recruited and mentored eight engineers and one manager, with responsibility for performance reviews, compensation decisions, delivery standards, staffing, architecture, and roadmap execution.
- Directed a distributed Agile organization through production launch for four large customers, including the U.S. Departments of State and Education.
- Designed forecasting, machine-learning, optimization, and scheduling capabilities integrated with Amazon Connect, Genesys, and Talkdesk.

Sr. Director, Data Analytics & Operations | TriMark USA | Mansfield, MA | Aug 2021 to May 2022

Built the cloud data organization and operating model supporting enterprise analytics across a \$1B+ company and 11 divisions.

- Upskilled an initial four-person analytics team and expanded it into a seven-member AWS engineering organization with additional contractors.

- Aligned executives around more than 50 use cases and delivered the first enterprise data lake within ten months, reducing reporting from roughly two months to one day.
- Owned staffing, performance reviews, architecture, delivery planning, risk, change management, executive reporting, and cross-functional priorities.

Engineering Lead, Cloud Data Privacy & AI Governance | MassMutual | Boston, MA | Nov 2020 to Aug 2021

Led architecture and engineering across privacy, governance, Responsible AI, data, and cloud initiatives in a regulated enterprise.

- Guided five initiatives spanning policy automation, model transparency, bias, drift, monitoring, human oversight, auditability, data governance, and governance APIs.
- Partnered with Enterprise Architecture and SRE on AWS, security, IAM, CI/CD, event-driven, container, and infrastructure-as-code standards.
- Delivered a cloud-native privacy platform within eight months that automated approximately 90% of a formerly manual workflow.

VP of AI/ML | NorthBay Solutions | Boston, MA | Jan 2019 to Nov 2020 | (Role ended in company-wide COVID layoff).

Recruited from AWS by the CEO to build an enterprise AI/ML consulting practice and reusable cloud delivery capabilities.

- Led an 18-person distributed engineering team delivering the Data Lake Accelerator while remaining hands-on in initial product architecture and implementation.
- Advised enterprise clients on AI, ML, data, platform strategy, architecture, and production delivery across concurrent consulting engagements.
- Combined executive advisory, staffing, reusable standards, cloud architectures, and production systems for enterprise customers.

Sr. Specialist Solutions Architect, AI/ML & Big Data | Amazon Web Services | Boston, MA | May 2015 to Dec 2018

Helped establish and scale AWS AI/ML and data specialist capability while advising enterprise customers on production adoption.

- Advised more than 200 enterprises across North America on AI/ML, data platforms, architecture, governance, cloud adoption, and business outcomes.
- Founded and led technical communities totaling 180+ specialists, created 170+ enablement assets, and established standards and reusable client practices.
- Delivered 36 invited presentations, trained 180+ architects, and earned AWS ML Elite Gold and Platinum recognition.

EARLIER RELEVANT EXPERIENCE

- Designed and shipped commercial scientific software for life-sciences customers, translating complex research workflows into maintainable product capabilities.
- Built cloud and device-software systems for an FDA-cleared connected product, including an AWS/Cassandra production platform.

EDUCATION & RECOGNITION

Ph.D., Molecular Biology / Microbiology & Immunology, University of Western Ontario. AWS ML Elite Gold and Platinum; AWS re:Invent Rockstar; named inventor on multiple patents.

TECHNOLOGIES

Python • FastAPI • REST APIs • backend services • agentic AI • AI agents • multi-agent systems • RAG • embeddings • vector databases • PostgreSQL • PGVector • LangChain • LangGraph • MCP • prompt engineering • AI evaluation • LLM-as-judge • failure analysis • guardrails • observability • reliability • security • data workflows • data pipelines • ETL • SQL • AWS • Amazon Bedrock • SageMaker • Lambda • Docker • Kubernetes • AWS CDK • CI/CD • Git • Linux • testing • technical documentation • cloud architecture • inference workflows • developer tooling • Responsible AI