

Keith Steward, Ph.D.

Data & AI Engineering Manager | Cloud Platforms

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

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

PROFILE

Player-coach data and AI engineering leader who builds governed cloud platforms, production pipelines, retrieval systems, and engineering teams that convert fragmented information into dependable business capability. Combines AWS, Snowflake, Spark, data architecture, RAG, vector retrieval, MLOps/LLMOps, security, governance, and hands-on Python with accountable team leadership. Built an 11-engineer AI organization and a production platform supporting 70+ automations across more than 20 workflows with 99.9% availability. Delivers scalable data foundations, clear standards, and measurable operating outcomes.

SELECTED CAREER HIGHLIGHTS

- Built and led an 11-engineer AI organization while owning architecture, hiring, mentoring, roadmaps, delivery, governance, cloud operations, and a \$900,000 budget.
- Personally built the AssistX core and first 80% of its initial application with Python, FastAPI, agents, RAG, PostgreSQL/PGVector, APIs, Docker, AWS, and CDK.
- Delivered 70+ automations across more than 20 workflows with 99.9% availability, implementing evaluation, observability, guardrails, model routing, cost controls, CI/CD, and incident response.
- Led enterprise data-product strategy across 11 businesses and more than 50 prioritized use cases, delivering an AWS data lake that reduced cross-ERP reporting from roughly two months to one day.
- Built and led data, analytics, AI/ML, cloud, and software teams as large as 25, establishing architecture standards, data governance, security controls, and production practices.
- Advised more than 200 AWS enterprise customers and built reusable production reference architectures across Snowflake, Spark, Redshift, data lakes, streaming, APIs, analytics, and machine learning.

PROFESSIONAL EXPERIENCE

Technical Co-Founder | Intelligent DataWorks | Boston, MA | Aug 2023 to Oct 2026

Built and operated a production AI company while personally owning strategy, architecture, implementation, delivery, governance, and cloud 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.
- Completed planned transition from operating CEO responsibilities to pursue a long-term 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

- Established NorthBay's AI/ML practice and built its first production Data Lake Accelerator hands-on before leading an 18-person distributed delivery team.
- Led regulated privacy and Responsible AI engineering at MassMutual, joining data governance, auditability, model monitoring, and secure production-ready cloud-platform delivery across enterprise programs.

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

data engineering • data architecture • AWS • Amazon Bedrock • SageMaker • Snowflake • Spark • Redshift • data lakes • ETL • data pipelines • Python • SQL • FastAPI • REST APIs • RAG • LangChain • LangGraph • MCP • PostgreSQL • PGVector • vector databases • embeddings • MLOps • LLMOps • AI evaluation • observability • data governance • AI governance • Responsible AI • security • privacy • CI/CD • AWS CDK • Docker • Kubernetes • microservices • event-driven systems • cost optimization • technical standards