

Keith Steward, Ph.D.

Principal Agentic AI Architect | Enterprise AI Delivery

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

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

PROFILE

Hands-on AI architect who takes agentic and conversational systems from discovery and rapid prototype through integration, production deployment, evaluation, and continuous optimization. Combines Python, agents, RAG, APIs, data pipelines, distributed cloud architecture, security, privacy, governance, and executive advisory with direct ownership of measurable outcomes. Personally built the core and first 80% of a production AI platform supporting 70+ automations across more than 20 workflows with 99.9% availability. Turns ambiguous enterprise needs into scalable architectures and durable operating capability.

SELECTED CAREER HIGHLIGHTS

- Built the AssistX core and first 80% of its initial application hands-on with Python, FastAPI, LLMs, agents, RAG, PostgreSQL/PGVector, REST APIs, Docker, AWS, and CDK.
- Translated more than 20 business workflows into 70+ production AI capabilities, accelerating selected work by up to 50x while sustaining 99.9% availability and supporting 150+ early users.
- Implemented structural and LLM-as-judge evaluation, prompt regression, human review, guardrails, prompt-injection screening, observability, auditability, model routing, release controls, and incident response.
- Advised more than 200 enterprise customers on AI/ML, data platforms, distributed cloud architectures, security, governance, and production adoption; created 170+ reusable technical assets.
- Delivered API-driven and event-oriented platforms across regulated environments, including a privacy-rights system that automated approximately 90% of a manual workflow.
- Combined product discovery, rapid prototyping, integration architecture, stakeholder workshops, demonstrations, production operations, and continuous improvement across enterprise AI initiatives.

PROFESSIONAL EXPERIENCE

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

Built and operated a production AI platform while personally owning product strategy, architecture, implementation, evaluation, delivery, security, 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 and Talkdesk in secure enterprise environments.

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

- Led an 18-person distributed engineering team while remaining hands-on in architecture and implementation of a reusable AWS Data Lake Accelerator.
- Led cloud and device-software engineering 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

Agentic AI • AI agents • conversational AI • NLP • virtual assistants • copilots • Python • FastAPI • REST APIs • JSON • event-driven systems • data pipelines • RAG • LangChain • LangGraph • MCP • LLMs • prompt engineering • AI evaluation • LLM-as-judge • prompt regression • human-in-the-loop review • guardrails • observability • monitoring • AWS • Amazon Bedrock • SageMaker • Google Vertex AI • distributed systems • PostgreSQL • PGVector • vector databases • Docker • Kubernetes • AWS CDK • CI/CD • security • privacy • data governance • Responsible AI • AI governance • HIPAA • executive communication • rapid prototyping