

Keith Steward, Ph.D.

Lead Agent Architect | Knowledge AI & Responsible Platforms

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

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

PROFILE

Hands-on AI platform architect and player-coach who builds high-autonomy agentic systems with robust knowledge, safety, governance, monitoring, and enterprise integration. Combines foundation-model selection, RAG, MCP, multi-agent orchestration, memory, evaluation, observability, security, Responsible AI, and cloud operations with mentoring and cross-functional leadership. Personally built the core and first 80% of a production platform, then led 11 engineers supporting 70+ automations across more than 20 workflows with 99.9% availability. Turns governance requirements into scalable working systems.

CORE TECHNOLOGIES

Enterprise AI Platforms • AI Governance • Responsible AI • IAM • SaaS Operations • ChatGPT • Claude • LLMs • AI Agents • RAG • Python • FastAPI • APIs • AWS • Observability

SELECTED CAREER HIGHLIGHTS

- Built and operates a production AI platform supporting 70+ automations across more than 20 workflows with 99.9% availability.
- Established model governance, evaluation, guardrails, human oversight, security, observability, model routing, and cost-control practices.
- Advised 200+ AWS enterprise customers and created reusable reference architectures, implementation guidance, workshops, and technical enablement.
- Led enterprise privacy, IAM, Responsible AI, auditability, and model-risk controls in financial services.
- Owned platform budgets up to \$900,000, cloud operations, vendor evaluation, adoption feedback, releases, and measurable workflow outcomes.
- Established reusable agent architecture patterns for reasoning, planning, tool use, retrieval, memory, orchestration, evaluation, guardrails, monitoring, and human oversight. Connected those patterns to model-provider controls, enterprise APIs, structured and unstructured knowledge, security boundaries, auditability, deployment automation, cost management, and incident response. Mentored engineers through architecture reviews and production troubleshooting while helping business, product, governance, and platform teams adopt agentic capabilities with shared standards and measurable quality gates across multiple production use cases and distributed teams at scale. Standardized Agentic AI and foundation models for dependable reuse across enterprise platforms.

PROFESSIONAL EXPERIENCE

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

Built and operated a production AI platform while personally owning architecture, implementation, engineering standards, evaluation, delivery, governance, 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.
- Owned 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

- Led enterprise Responsible AI, privacy, IAM, auditability, and MLOps governance for MassMutual.
- Created reusable AWS AI/ML and data accelerators and enabled global specialist communities totaling 180+ professionals.

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

enterprise AI platforms • ChatGPT • Claude • Generative AI • LLMs • AI agents • RAG • LangChain • LangGraph • MCP • Python • FastAPI • REST APIs • workflow automation • SaaS operations • IAM • RBAC • AI governance • Responsible AI • security • privacy • auditability • guardrails • human-in-the-loop • AI evaluation • LLM-as-judge • observability • monitoring • cost controls • model routing • AWS • Amazon Bedrock • SageMaker • PostgreSQL • PGVector • Docker • Kubernetes • AWS CDK • CI/CD • incident response • vendor evaluation • technical enablement