

Keith Steward, Ph.D.

VP Engineering Leader | Healthcare AI & Scalable Platforms

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

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

PROFILE

Engineering executive who builds secure healthcare and enterprise AI platforms while scaling hands-on development teams and operating practices. Combines technical strategy, cloud-based systems, data platforms, AI automation, HIPAA-regulated software, security, health data privacy, and reliable production delivery. Personally built the core and first 80% of an AI platform, then recruited and led 11 engineers around 70+ automations across more than 20 workflows with 99.9% availability. Moves from roadmap and build-versus-buy choices through architecture, implementation, hiring, governance, and measurable operations improvement.

SELECTED CAREER HIGHLIGHTS

- Personally built the AssistX core and first 80% of its initial application with Python, FastAPI, LLMs, RAG, agents, PostgreSQL/PGVector, APIs, Docker, AWS, and CDK, then recruited and led 11 engineers through production delivery.
- Delivered the first commercial release within eight months; the platform supports 70+ automations across more than 20 backend workflows, accelerates selected work by up to 50x, and sustains 99.9% availability.
- Defined technical strategy, architecture, engineering roadmaps, build-versus-buy choices, team structure, budgets, vendors, security, privacy, release controls, observability, incident response, and production operations across startup and enterprise environments.
- Led HIPAA-regulated and FDA-cleared healthcare technology, including a fault-tolerant AWS/Cassandra platform, APIs, device integrations, monitoring, analytics, high availability, and a five-person engineering team with a sub-team manager.
- At AWS, advised more than 200 enterprises—including healthcare organizations—on AI/ML, data, cloud architecture, security, reliability, and adoption while developing global technical communities totaling 180+ specialists.
- Built organizations as large as 25 and led teams of 7, 11, 24, and 25 through hiring, mentoring, manager development, architecture, product delivery, operational metrics, and cross-functional execution in ambiguous, high-growth settings.

PROFESSIONAL EXPERIENCE

Technical Co-Founder & CEO | 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 CambridgeSoft's 25-person engineering organization delivering scientific SaaS for pharmaceutical and biotechnology customers, combining commercial products, enterprise data, cloud systems, and team development.
- Built an award-winning, HIPAA-regulated and FDA-cleared digital-dentistry platform with fault-tolerant AWS/Cassandra infrastructure, APIs, security, monitoring, analytics, and no significant production downtime.
- Holds a Ph.D. in Molecular Biology and Microbiology & Immunology and developed commercial bioinformatics and knowledge-discovery software earlier in his career.

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

engineering strategy • technology roadmaps • organization design • hiring • manager development • budget ownership • vendor strategy • build-versus-buy • healthcare technology • HIPAA • FDA compliance • security • privacy • Generative AI • LLMs • RAG • AI agents • workflow automation • AI evaluation • Responsible AI • Python • FastAPI • REST APIs • microservices • distributed systems • AWS • SageMaker • data platforms • data pipelines • PostgreSQL • PGVector • Cassandra • LangChain • LangGraph • MCP • Docker • Kubernetes • AWS CDK • CI/CD • observability • reliability • incident response