

Keith Steward, Ph.D.

Generative AI Engineering Leader | Distinguished AI Platform Architect

Boston, MA | 508-318-8035 | keith.steward@gmail.com | LinkedIn | IntelligentDataWorks.ai | IDW LinkedIn

PROFILE

Hands-on AI platform architect and engineering leader who turns GenAI, data, and cloud architecture into production systems that teams can trust. Built AssistX, a production-grade LLM/RAG/agent platform with 70+ AI automations across 20+ workflows, while recruiting and leading 11 AI engineers. Previously advised 200+ AWS enterprise customers and developed 180+ AWS specialists. Brings the technical depth, evaluation discipline, and operating judgment to establish a durable GenAI foundation and grow a lean, high-bar team.

SELECTED QUALIFICATIONS

- **GenAI platform architecture:** LLMs, RAG, AI agents, context layers, vector search, model-serving tradeoffs, evaluation pipelines, guardrails, and workflow orchestration.
- **Hands-on production delivery:** Python, FastAPI, PostgreSQL/PGVector, LangChain, LangGraph, AWS, Docker, CI/CD, data pipelines, APIs, cloud operations, and system integration.
- **Technical leadership:** architecture decisions, code review, hiring bar, engineering standards, agile delivery, mentoring, team growth, and senior stakeholder communication.
- **Enterprise scale:** AWS AI/ML and big-data architecture guidance for 200+ customers, production data platforms, governance systems, SaaS products, and distributed engineering teams.
- **Responsible AI and reliability:** LLM-as-judge evals, prompt controls, input validation, observability, auditability, security, and human oversight.

SELECTED CAREER HIGHLIGHTS

- Built and helped commercialize six software products, including AssistX, Ingenuity Pathway Analysis, BioWizard, CambridgeSoft Biological Notebook, Lythos cloud software, and AWS data/ML accelerators.
- Delivered AssistX's first commercial release in eight months, supporting hundreds of user signups, 70+ AI automations, and selected workflows running up to 50x faster than manual processing.
- Created AWS AI/ML and Big Data technical communities, trained 180+ solution architects, produced 170+ enablement assets, and earned AWS ML Elite Gold and Platinum recognition.

PROFESSIONAL EXPERIENCE

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

Lead AI product strategy, architecture, hands-on platform development, engineering execution, and operations for an applied GenAI SaaS platform.

- Architected and built the AssistX core platform and the first 80% of the HR assistant, integrating LLMs, RAG, AI agents, LangChain, LangGraph, PostgreSQL/PGVector, Python, FastAPI, AWS, Docker, and MLOps patterns.
- Delivered a production-grade GenAI assistant with 70+ AI automations across more than 20 workflows, including labor-law RAG, resume evaluation, HR-policy Q&A, voice navigation, and workflow orchestration.
- Implemented LLM-as-judge evals, synthetic test data, prompt controls, user-input validation, prompt-injection defenses, observability, and quality reviews.
- Recruited, trained, and technically led an 11-engineer Agile AI team selected from 400+ applicants while staying directly involved in architecture, code, cloud operations, security, releases, and system administration.
- Initiated strategic acquisition discussions regarding IDW's AI platform as part of a planned transition to a senior AI/engineering leadership role.

Senior Director of Engineering | SuccessKPI | Remote | Jun 2022 to Aug 2023

Built and led engineering for a cloud-native, AI-powered Workforce Management product within an enterprise contact-center analytics SaaS platform.

- Recruited and mentored a distributed engineering team delivering forecasting, scheduling, optimization, analytics, automated communication, and platform services on AWS.
- Translated product requirements into architecture, algorithms, APIs, cloud deployment, testing, CI/CD, and two-week Agile releases integrated with Amazon Connect, Genesys, and Talkdesk.

- Designed ML and operations-research capabilities using Python, Aurora PostgreSQL, AWS services, and linear-programming optimization to convert demand and staffing constraints into schedules.
- Delivered production access for large customers including the U.S. Department of State, Avon, De'Longhi, and the Department of Education within 15 months.

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

Owned enterprise data modernization for a \$1B+ organization integrating 11 acquired businesses and heterogeneous ERP systems.

- Defined the AWS data strategy, roadmap, and architecture with Finance, Marketing, Supply Chain, and technology leaders, prioritizing 50+ analytics use cases and KPIs.
- Upskilled an initial analytics team and grew it to seven AWS engineers plus contractors while managing planning, budget, risk, delivery, and executive communication.
- Delivered the first AWS data lake within seven months using Glue, S3, PySpark, Data Catalog, CDK, and automated pipelines, reducing cross-ERP reporting from roughly two months to one day.

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

Led architecture and implementation across Data Privacy, Responsible AI, and AI Governance initiatives in a regulated financial-services environment.

- Converted privacy, governance, and policy requirements into secure AWS architecture, APIs, event-driven integration, CI/CD, serverless patterns, and production delivery.
- Delivered the first cloud-native privacy-rights management platform within eight months, automating approximately 90% of a manual CCPA workflow.

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

Recruited out of AWS by the NorthBay Solutions CEO to establish a new AI/ML practice and reusable data-and-ML accelerator.

- Built the initial Data Lake Accelerator hands-on with Python, AWS, APIs, and full-stack components, then directed a remote team toward customer demonstration and delivery.
- Advised enterprise customers including Dun & Bradstreet, Harvard Business School, 1091 Media, and Invesco on AI strategy, data platforms, knowledge graphs, solution architecture, and production implementation.
- Delivered an AWS Neptune knowledge graph integrating complex enterprise hierarchies and expanded reusable cloud, data, and ML architectures for customer engagements.

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

Advised enterprises and AWS field teams on production machine learning, distributed analytics, data platforms, and cloud adoption.

- Advised 200+ enterprise customers on AI/ML strategy, SageMaker, TensorFlow, data lakes, EMR/Spark, security, distributed systems, and production cloud architecture.
- Created and led global Big Data and AI/ML technical communities, developing 180+ specialists through standards, mentoring, workshops, hackathons, and reusable assets.
- Created 170+ technical assets, delivered 50+ presentations, built a SageMaker/TensorFlow hackathon for Morgan Stanley, and earned AWS ML Elite Gold and Platinum recognition.

EARLIER RELEVANT EXPERIENCE

- Director of Software at Dimensional Photonics: led a cloud-connected dental-scanner software team through FDA-cleared product launch, AWS/Cassandra platform delivery, APIs, analytics, and high-availability operations.
- Earlier AI work includes BioWizard, a 400-rule inferencing system for genomic drug-target evaluation, and original development work on Ingenuity Pathway Analysis, a globally adopted commercial bioinformatics product.

EDUCATION & RECOGNITION

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

TECHNOLOGIES

LLMs • OpenAI API • Anthropic Claude • AWS Bedrock • RAG • AI agents • multi-agent systems • LangChain • LangGraph • prompt engineering • LLM-as-judge evals • vector databases • PostgreSQL • PGVector • Python • FastAPI • REST APIs • AWS • SageMaker • EMR • Spark • PySpark • Glue • Docker • Kubernetes • CI/CD • Linux • observability • security