

Keith Steward, Ph.D.

Principal AI Engineer | Applied AI Platform Architect

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

PROFILE

Principal-level applied AI engineer and hands-on platform builder who turns LLM capabilities into reliable, reusable product foundations. Personally architected the AssistX core and approximately 75% of its initial application, supporting 70+ AI capabilities across more than 20 workflows with 99.9% availability. Combines Python, RAG, agentic workflows, prompt engineering, evaluation harnesses, APIs, reusable libraries, observability, and cloud-native delivery. Brings rapid prototyping, production hardening, architecture judgment, and consultative technical influence across product and engineering teams.

AREAS OF EXPERTISE

Applied AI Systems: Python · LLMs · prompt engineering · RAG · agentic workflows

Evaluation & Reliability: test harnesses · LLM-as-judge · observability · security · guardrails

Reusable Engineering: APIs · shared libraries · model clients · workflow orchestration · cloud delivery

Technical Influence: architecture tradeoffs · rapid prototypes · production hardening · product partnership

SELECTED CAREER HIGHLIGHTS

- Delivered the initial production AssistX platform in approximately eight months while balancing rapid delivery with enterprise-grade scalability, security, maintainability, and operational excellence.
- Advised more than 200 AWS enterprise customers on AI/ML, data, analytics, cloud adoption, and production architecture; earned AWS ML Elite Gold and Platinum recognition.
- Founded and led AWS technical communities that prepared 180+ global Solution Architects and produced 170+ enablement assets for emerging data and AI services.
- Delivered TriMark's first AWS data lake across 11 divisional ERP systems, reducing cross-enterprise reporting from roughly two months to one day.
- Delivered MassMutual's first cloud-native privacy-rights platform in eight months, automating approximately 90% of a formerly manual regulated workflow.

PROFESSIONAL EXPERIENCE

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

Owned company strategy, product direction, architecture, engineering, hiring, budget, delivery, and platform operations for a reusable enterprise generative AI platform and HR assistant.

- Personally architected and built the AssistX core and approximately 75% of the initial HR application, then directed an 11-engineer Agile team through production releases on AWS.
- Designed a composable platform using Python, FastAPI, PostgreSQL/PGVector, LLM APIs, RAG, LangChain, LangGraph, agents, REST APIs, MCP servers and clients, workflow orchestration, and AWS CDK.
- Established model flexibility, prompt version control, LLM-as-judge and structural evaluations, synthetic-data testing, prompt-injection screening, human review, observability, and operational controls.
- Translated customer discovery and domain-expert input into a product roadmap; launched to 150+ early adopters and pilot users within approximately eight months.
- Directed implementation of LangChain, LangGraph, PostgreSQL/PGVector, AWS cloud services, multi-agent architectures, and retrieval-augmented generation pipelines.

Sr. Director of Engineering | SuccessKPI | Jun 2022 to Aug 2023

Led architecture, engineering, and delivery of a new Workforce Management capability within an AI-powered, AWS-based contact-center analytics SaaS platform.

- Recruited, coached, and managed eight engineers and one manager across a distributed team; owned performance reviews, compensation input, delivery planning, architecture, and stakeholder updates.
- Directed forecasting, scheduling, optimization, analytics, and platform integrations using Python, linear programming, Amazon S3, Aurora PostgreSQL, and enterprise CCaaS platforms.
- Delivered the initial production platform within 14 months for four customers, including two U.S. government departments, while operating within FedRAMP-informed requirements.

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

Modernized data and analytics for a \$1B+ enterprise with 11 acquired divisions and heterogeneous ERP systems.

- Defined the enterprise data strategy, secured executive support, prioritized 50+ business use cases, and grew an initial four-person analytics team into a seven-member AWS engineering organization.
- Delivered a 3 TB AWS data lake within 10 months using Glue, PySpark, S3, JDBC, and AWS CDK to ingest, catalog, consolidate, and normalize data across all 11 ERPs.
- Enabled consolidated analytics for finance, supply chain, sales, marketing, and operations while improving reporting turnaround by approximately 60x.

Engineering Lead, Cloud Data Privacy & AI Governance | *MassMutual* | Nov 2020 to Aug 2021

Established cloud-native engineering patterns for enterprise privacy, AI governance, and responsible AI initiatives in insurance and financial services.

- Directed architecture and hands-on implementation for privacy-rights management while guiding enterprise data catalog, adaptive-policy, permissible-use, and governance API initiatives.
- Partnered with enterprise architecture, security, and site reliability teams on AWS, IAM, serverless, integration, and CI/CD patterns that translated policy into practical controls.
- Advised AI Governance & Ethics leaders on model transparency, bias detection, drift monitoring, auditability, and lifecycle oversight using Amazon SageMaker.

Vice President, AI/ML | *NorthBay Solutions* | Jan 2019 to Nov 2020 (Role ended in company-wide COVID layoff).

Recruited from AWS by the CEO to establish the firm's enterprise AI consulting practice and expand its cloud, data, and machine-learning portfolio.

- Won CEO approval for a reusable Data Lake Accelerator, built the initial full-stack system hands-on, and led 18 offshore engineers through product completion and customer demonstrations.
- Advised enterprise clients including Dun & Bradstreet, Harvard Business School, 1091 Media, and Invesco on data strategy, AI/ML architecture, and AWS production delivery.
- Combined executive advisory, product management, reusable architectures, and hands-on implementation to expand repeatable enterprise AI and analytics offerings.

Sr. Specialist Solutions Architect, AI/ML and Big Data/Analytics | *Amazon Web Services* | May 2015 to Dec 2018

Served as AWS's first Senior Specialist Solutions Architect for AI/ML after leading Big Data and Analytics specialization during early enterprise adoption of EMR and SageMaker.

- Advised 200+ customers across financial services, healthcare, telecommunications, manufacturing, and technology on AI/ML, data lakes, analytics, governance, and production cloud architectures.
- Founded and led 100-member Big Data/Analytics and 80-member AI/ML technical communities; developed enablement, demonstrations, workshops, and onboarding for new services.
- Delivered 36 invited presentations, contributed to two AWS specialty certification exams, tested emerging services, and brought customer feedback to product teams.

PATENTS, RECOGNITION & EDUCATION

Patents: U.S. Patents 8,793,073 and 8,489,334; BioWizard Intelligent Agent Architecture patent application.

Recognition: AWS ML Elite Gold · AWS ML Elite Platinum · AWS re:Invent Rockstar · AWS Certification SME.

Education: Ph.D., Molecular Biology / Microbiology & Immunology; Hon. B.Sc., Microbiology & Immunology — University of Western Ontario. Canadian Medical Research Council Post-Doctoral Research Fellowship.