

Keith Steward, Ph.D.

Chief Technology Officer | Chief AI Officer | AI Platform & Product Executive | Data & Cloud Leader

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

<https://www.linkedin.com/in/keithsteward/> | <https://IntelligentDataWorks.ai> | <https://bit.ly/3Th0CFX>

EXECUTIVE PROFILE

Executive AI, machine learning, data, and cloud technology leader with 25+ years creating production AI/ML, analytics, SaaS, and enterprise software systems. Operated as founder/CEO, VP AI, senior engineering director, data analytics executive, AI governance technical lead, and AWS senior specialist across AI/ML and Big Data. Solves business problems where unstructured data, manual workflows, product capability, and enterprise adoption need to become scalable AI systems.

Combines hands-on LLM/NLP/RAG architecture, full ML lifecycle judgment, cloud deployment depth, product leadership, team building, and C-suite advisory experience. Has built AI products and practices from zero, advised 200+ enterprise AWS customers, directed distributed engineering teams, and delivered production systems across SaaS, HR/HCM, contact-center analytics, financial services, enterprise data platforms, and biotechnology software.

Career Note: Following four years at AWS, I accepted three successive executive leadership opportunities to build new AI, Data Analytics, and governance capabilities in organizations undergoing significant transformation. In 2023 co-founded *Intelligent DataWorks* inc. to commercialize an enterprise GenAI + Business Data platform. Now seeking a long-term executive leadership opportunity following a planned strategic sale of that business.

EXECUTIVE VALUE PROPOSITION

Executive leader combining AI strategy, deep technical expertise, product leadership, enterprise transformation, startup execution, commercialization, organization building, AI governance, and hands-on production delivery.

AI SKILLS / TECHNOLOGIES

AI Leadership: Enterprise AI strategy, applied AI productization, AI platform ownership, AI governance, responsible AI, executive stakeholder advisory, product-engineering alignment, team building, roadmap sequencing, adoption planning.

Generative AI: LLMs, AI agents, agentic workflows, LangChain, LangGraph, OpenAI APIs, Anthropic Claude, RAG, embeddings, vector search, multi-agent systems, prompt/workflow design, evaluation, cost controls, hallucination minimization.

Machine Learning: NLP, document intelligence, unstructured data, forecasting, predictive analytics, optimization, linear programming, model explainability, bias/drift/transparency concepts, MLOps, TensorFlow, SageMaker, SparkML, observability.

Cloud / Data / Engineering: AWS SageMaker, Bedrock, EMR, Glue, Athena, S3, Lambda, Step Functions, API Gateway, Aurora, CDK, CloudFormation, Kubernetes, Docker, Python, SQL, PostgreSQL/PGVector, FastAPI, Flask, Pandas, PySpark, Git, CI/CD, REST APIs, distributed systems, SaaS architecture.

SELECTED CAREER HIGHLIGHTS

- Founded and directed an AI startup and 11-engineer AI organization.
- Created *AssistX*, a production GenAI assistant platform for 70+ automations across 20+ HR workflows.
- Supported AI and analytics adoption across 200+ AWS enterprise customers.
- Created AWS AI/ML technical communities and trained 180+ global solution architects.
- Delivered TriMark's first AWS data lake, reducing cross-ERP reporting from roughly 2 months to 1 day.
- Technically directed Data Privacy and AI Governance cloud engineering at MassMutual.
- Implemented commercial software products including the globally adopted Ingenuity Pathway Analysis platform.
- Earned AWS ML Elite Gold and Platinum recognition; named inventor on multiple patents.

PROFESSIONAL EXPERIENCE

Technical Co-Founder & CEO | *Intelligent DataWorks inc* | Aug 2023 to Present

Founded an applied AI SaaS company to commercialize an enterprise GenAI + business data platform. Owns AI strategy, product vision, platform architecture, engineering execution, workflow design, security posture, customer validation, and strategic transition planning.

- Architected *AssistX*, a production-grade generative AI assistant platform using LLMs, AI agents, RAG, PostgreSQL/PGVector, LangChain, LangGraph, APIs, workflow orchestration, and AWS cloud services.
- Created reusable AI platform foundations for agentic workflows, RAG grounding, vector search, evaluation loops, observability, human-in-the-loop quality, secure data handling, and production adoption.
- Produced 70+ AI automations across 20+ HR and business workflows, improving selected workflows by up to 50x versus manual execution.
- Directed an 11-engineer Agile AI development team while staying hands-on in product architecture, roadmap execution, cloud engineering, backend systems, AI orchestration, and platform operations.
- Delivered the initial production release within 8 months and converted emerging LLM capability into a durable SaaS platform.
- Established AI product strategy, roadmap sequencing, stakeholder narratives, adoption plans, commercialization assets, and business-value messaging for enterprise AI workflows.
- Applied Python, PostgreSQL/PGVector, FastAPI, AWS, LLM APIs, RAG, agentic workflows, and evaluation practices to production AI workflows.
- Positioned the business for strategic transition after validating product direction, platform architecture, and enterprise workflow value.

Sr. Engineering Director, Workforce Management Service | *SuccessKPI* | Jun 2022 to Aug 2023

Directed engineering, architecture, and delivery for the Workforce Management component of an AI-powered contact-center analytics SaaS platform. Converted product requirements into cloud architecture, forecasting algorithms, optimization workflows, integrations, and release plans.

- Recruited and directed a 7-member distributed engineering team building AWS-based forecasting, scheduling, visualization, and automated communication capabilities.
- Designed AI-driven forecasting, optimization, workflow automation, API integration, analytics, and product delivery patterns for enterprise SaaS and customer-facing operations.
- Delivered the first working WFM release within 6 months for two flagship enterprise customers and advocated early generative AI adoption across the platform.

Sr. Director, Data Analytics & Operations | *TriMark USA* | Aug 2021 to May 2022

Recruited to build the company's first enterprise data lake initiative across 11 ERP systems and establish enterprise analytics capability. Partnered with marketing, supply chain, finance, and executive stakeholders to prioritize analytics use cases and KPI delivery.

- Secured support for an AWS data lake strategy enabling data engineering, analytics, and future machine learning capabilities across a \$1B+ enterprise.
- Delivered TriMark's first AWS data lake within 7 months, automating nightly ingest, cataloging, consolidation, and normalization across 11 ERPs.
- Improved cross-ERP reporting cycles from roughly 2 months of manual work to 1 day, a 60x improvement in insight delivery.

Technical Lead for Cloud, Data Privacy and AI Governance | *MassMutual* | Nov 2020 to Aug 2021

Established AWS architecture and hands-on implementation for Data and AI Governance capabilities, including privacy rights management, enterprise data catalog, adaptive policy framework, governance APIs, and AI governance initiatives.

- Coordinated product managers, enterprise architecture, governance, site reliability, and developers to convert requirements into secure cloud-native architecture, sprint plans, CI/CD, and production delivery.
- Provided MLOps guidance for model transparency, bias, drift detection, auditability, monitoring, and lifecycle management using Amazon SageMaker.
- Delivered a first cloud-native privacy rights platform in 8 months, automating roughly 90% of a formerly manual workflow.

VP of Machine Learning & AI | NorthBay Solutions | Jan 2019 to Nov 2020

Recruited out of AWS by the NorthBay Solutions CEO to head up a new AI/ML practice for enterprise customers and AWS-based product acceleration. Combined customer advisory, hands-on engineering, solution architecture, and executive delivery.

- Generated CEO approval for an AWS Data Lake + ML Accelerator SaaS initiative, built the initial working system hands-on, and directed remote engineering toward customer demonstrations.
- Delivered AI/ML and analytics solutions for Dun & Bradstreet, Harvard Business School, Invesco, and other enterprise customers.
- Departed after the 2020 COVID pandemic wiped out NorthBay's revenues and triggered a large company layoff that included this role.

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

Served as an AWS senior specialist across AI/ML, Big Data, analytics, data lakes, machine learning, and enterprise cloud adoption. Advised senior customer stakeholders and enabled AWS field teams with technical strategy, architecture, demos, and training.

- Created and grew AWS internal AI/ML and analytics technical communities, enabling 180+ global solution architects and producing 170+ technical assets.
- Advised 200+ enterprise customers on AI strategy, ML, data lake architecture, SaaS adoption, governance, security-aware architecture, SageMaker, TensorFlow, Glue, Athena, S3, Lambda, and production patterns.
- Presented 50+ formal AI/ML and analytics sessions and earned AWS ML Elite Gold and Platinum recognition.

PATENTS / PUBLICATIONS / SPEAKING

- United States Patent No. US 8,793,073.
- United States Patent No. US 8,489,334.
- BioWizard Intelligent Agent Architecture patent application (1998) for genomic knowledge discovery.
- Delivered 36 formal presentations on Machine Learning, Deep Learning, and Artificial Intelligence at large AWS and non-AWS/industry events.
- AWS ML Elite Gold and AWS ML Elite Platinum awards.
- AWS re:Invent Rockstar and AWS Certification SME recognition.

EDUCATION

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