

Keith Steward, Ph.D.

AI Outcome Customer Engineer | Forward-Deployed AI Architecture | Cloud & Data Executive

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, cloud, data, and software technology leader with 25+ years creating production AI/ML, analytics, SaaS, and enterprise platforms. Operated as founder/CEO, VP AI, senior engineering director, data analytics executive, AI governance technical lead, and AWS senior specialist.

Solves ambiguous customer problems where enterprise architecture, technical debugging, code-level context, identity, compliance boundaries, and AI model deployments must become viable delivery. Combines hands-on engineering depth with executive communication and customer-facing technical authority.

Career Note: Following four years at AWS, 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 LEADERSHIP & CORE EXPERTISE

AI Strategy & Executive Leadership: Enterprise AI strategy, customer engineering, forward deployed engineering, technical authority, architectural foresight, adoption strategy, executive communication, product-engineering liaison.

Business Transformation & Delivery: Strategic account feasibility, pre-sales technical evaluation, post-sales execution, rapid activation, viable delivery, implementation strategy, customer success, business impact.

Generative AI: LLMs, AI agents, model deployments, RAG, embeddings, vector search, LangChain, LangGraph, OpenAI APIs, Anthropic Claude, evaluation, human-in-the-loop controls.

AI Engineering & Machine Learning: System design, Python, troubleshooting, code-level debugging, model explainability, MLOps, TensorFlow, SageMaker, SparkML, predictive analytics, optimization.

Enterprise Data & Cloud Platforms: AWS architecture, cloud infrastructure, APIs, connectors, identity, data pipelines, privacy, data residency, compliance boundaries, PostgreSQL/PGVector, Kubernetes, Docker, CI/CD.

SELECTED CAREER HIGHLIGHTS

- Former AWS AI/ML and Big Data senior specialist advising 200+ enterprise customers on AI strategy, architecture, model deployments, data lakes, and production adoption.
- Founded and directed an AI SaaS startup and 11-engineer AI organization building production GenAI assistant workflows.
- Created *AssistX*, a production GenAI assistant platform for 70+ automations across 20+ workflows.
- Built AI platform foundations spanning agents, RAG, vector search, model evaluation, observability, identity-aware workflow integration, privacy, and responsible AI controls.
- Directed regulated financial-services privacy, cloud, data, and AI governance delivery involving compliance boundaries and enterprise architecture.
- Delivered TriMark's first AWS data lake, reducing cross-ERP reporting from roughly two months to one day.
- Built and scaled engineering organizations delivering enterprise AI platforms, cloud-native software, analytics, integrations, and SaaS products.
- Delivered AI and bioinformatics products from concept through production, commercialization, and customer adoption.

PROFESSIONAL EXPERIENCE

Technical Co-Founder & CEO | *Intelligent DataWorks inc* | Boston, MA | Aug 2023 - 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.

- Operated as founding AI executive and technical delivery manager for a customer-facing AI SaaS business, validating AI outcome, project feasibility, FDE / Forward Deployed Engineering fit, pre-sales technical evaluation, post-sales execution, rapid activation, viable delivery, and long-term adoption.
- Architected *AssistX*, a production generative AI assistant platform using LLMs, AI agents, RAG, PostgreSQL/PGVector, LangChain, LangGraph, APIs, connectors, workflow orchestration, and AWS cloud services.
- Created reusable AI platform foundations for enterprise architecture, system design, models, model deployments, customer data pipelines, identity providers, enterprise content management / ECM workflows, data controls, compliance boundaries, and credible architectural solutions.
- Dove into Python, API, retrieval, prompt, data, and orchestration code-level context, reading code to solve novel technical problems, decipher ambiguity, diagnose bugs, test workarounds, clear execution roadblocks, and convert implementation issues into actionable product fixes.
- Directed an 11-engineer Agile AI team while staying hands-on in technical design, cloud infrastructure, backend systems, AI orchestration, technical debugging, systemic debugging, deployment quality, implementation strategy, technical authority, architectural foresight, and delivery strategy.
- Produced 70+ AI automations across 20+ business workflows, improving selected workflows by up to 50x and translating ambiguous customer needs into measurable business impact.
- Served as product and engineering liaison for roadmap prioritization, feature requests, stakeholder narratives, executive communication, customer success activation, technical go-to-market support, and cloud era adoption messaging.

Sr. Engineering Director, Workforce Management Service | *SuccessKPI* | Centreville, VA | Jun 2022 - 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 cloud infrastructure, forecasting, scheduling, visualization, enterprise integrations, and automated communication capabilities.
- Designed AI-driven forecasting, optimization, API integration, troubleshooting, customer-facing delivery, and enterprise SaaS implementation patterns across Amazon Connect, Genesys, Talkdesk, and CCaaS workflows.
- 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* | Mansfield, MA | Aug 2021 - 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 pipelines, analytics, governance, executive communication, and future machine learning capabilities across a \$1B+ enterprise.
- Delivered TriMark's first AWS data lake within 7 months, automating nightly ingest, cataloging, consolidation, normalization, and cross-system data integration across 11 ERPs.
- Improved cross-ERP reporting cycles from roughly two months of manual work to one day, a 60x improvement in customer-facing insight delivery.

Technical Lead for Cloud, Data Privacy and AI Governance | *MassMutual* | Boston, MA | Nov 2020 - 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, privacy governance, AI governance, security stakeholders, site reliability, and developers to convert requirements into secure cloud-native architecture and production delivery.

- Provided MLOps guidance for model transparency, bias, drift detection, auditability, monitoring, lifecycle management, data residency concerns, legal constraints, and compliance evidence using Amazon SageMaker.
- Delivered a cloud-native privacy rights platform in 8 months, automating roughly 90% of a formerly manual workflow while respecting identity, privacy, and compliance boundaries.

VP of Machine Learning & AI | NorthBay Solutions | Boston, MA | Jan 2019 - 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.

- Recruited out of AWS by the CEO to establish a new AI/ML practice for strategic accounts, cloud architecture, customer engineering, and enterprise AI 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, analytics, system design, enterprise architecture, and data 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 | Boston, MA | May 2015 - 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 for customer engineers and account teams.
- Advised 200+ enterprise customers on AI/ML model deployments, ML, deep learning, AI strategy, EMR/Spark, data lake architecture, SageMaker, TensorFlow, Glue, Athena, S3, Lambda, APIs, and production adoption patterns.
- Presented 36+ formal AI/ML and analytics sessions, translated deep technical issues into business impact, 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., Microbiology & Immunology, University of Western Ontario, Canada, 1993.
- Hon. B.Sc., Biology, University of Western Ontario, Canada, 1986.
- Canadian Medical Research Council Post-Doctoral Research Fellowship, 1993.

TECHNICAL SKILLS

Python, SQL, Java/C++ code-reading fluency, system design, troubleshooting, APIs, connectors, cloud infrastructure, AWS SageMaker, Bedrock, EMR, Glue, Athena, S3, Lambda, Step Functions, API Gateway, Aurora, CDK, CloudFormation, ECS/Fargate, Kubernetes, Docker, EventBridge, MSK/Kafka, Control Tower, PostgreSQL/PGVector, FastAPI, Flask, Pandas, PySpark, Git, CI/CD, REST APIs, SaaS architecture, enterprise integrations, identity, data pipelines, model deployments, LLMs, AI agents, RAG, vector search, MLOps, observability, privacy, compliance, and AI governance.