

Keith Steward, Ph.D.

VP, AI & Software Engineering | AI-Native Platform Executive

Boston, MA | 508-318-8035 | keith.steward@gmail.com | [LinkedIn](#) | [IntelligentDataWorks.ai](#) | <https://bit.ly/3Th0CFX>

EXECUTIVE PROFILE

AI and software engineering executive with 25+ years building teams, modernizing enterprise platforms, and converting emerging technology into production systems. Solves the organizational and architectural friction that prevents AI-assisted delivery, agentic applications, and legacy modernization from scaling safely. Built a production GenAI platform supporting 70+ automations across 20+ workflows, directed distributed teams up to 24, and advised 200+ AWS enterprise customers. Brings hands-on engineering depth, governance discipline, financial stewardship, and operating-model leadership to accelerate reliable delivery and measurable value.

EXECUTIVE LEADERSHIP & CORE EXPERTISE

- **AI-native engineering:** AI-assisted SDLC, agentic applications, developer enablement, evaluation, guardrails, reliability, and observability.
- **Software and platform modernization:** Legacy evolution, modular and API-first architecture, cloud-native systems, enterprise integration, and production operations.
- **Organization leadership:** Distributed teams, managers and staff, hiring, mentoring, engineering standards, Agile delivery, partners, and sourcing models.
- **Technology strategy and governance:** Roadmaps, architecture principles, responsible AI, investment priorities, build-buy-partner decisions, security, and risk.
- **Business and financial outcomes:** Budget planning, value-based prioritization, operating discipline, cloud economics, productivity, and measurable delivery results.

SELECTED CAREER HIGHLIGHTS

- Created AssistX, a production GenAI platform supporting 70+ automations across 20+ workflows, with up to 50x acceleration and projected savings above \$50K per customer.
- Built and led distributed engineering organizations up to 24 people while remaining hands-on in architecture, software delivery, cloud platforms, and AI systems.
- Advised 200+ AWS enterprise customers on AI/ML, analytics, data lakes, cloud architecture, and production adoption.
- Delivered TriMark's first data lake across 11 divisions, reducing cross-ERP reporting from about two months to one day—a 60x improvement.
- Established AWS technical communities of 80 and 100 specialists, trained 180+ architects, and produced 170+ reusable technical assets.

PROFESSIONAL EXPERIENCE

Technical Co-Founder & CEO | Intelligent DataWorks inc | Boston, MA | Aug 2023–Present

Own AI strategy, product and platform architecture, engineering execution, operating priorities, and enterprise-readiness for an applied-AI SaaS company.

- Architected and helped build AssistX, a production GenAI platform combining LLMs, AI agents, RAG, workflow orchestration, APIs, enterprise data, and AWS cloud services.
- Introduced AI-assisted engineering and code-generation practices that produced more than 50% of the platform code while maintaining human review, controlled prompts, input validation, and production guardrails.
- Recruited and directly led an 11-engineer Agile organization across AI, full-stack software, cloud, and DevOps while remaining hands-on in architecture, coding, integration, deployment, and operations.
- Delivered the first production release in eight months; supported 70+ automations across 20+ workflows, accelerating selected work up to 50x with projected savings above \$50K annually per customer.
- Initiated strategic acquisition discussions regarding IDW's AI platform as part of a planned transition to a senior AI/engineering leadership role.

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

Led architecture, engineering, and delivery for an AI-powered workforce-management SaaS product integrated with established contact-center platforms.

- Recruited and led a seven-member distributed engineering team through Agile sprints, standards, code quality, architecture decisions, release planning, and customer delivery.
- Delivered the initial production platform within six months, combining ML forecasting, operations-research scheduling, workflow automation, APIs, and AWS cloud architecture.
- Translated detailed product requirements into architecture, algorithms, delivery plans, and measurable product capabilities for major enterprise and government customers.

Sr. Director, Data Analytics & Operations | TriMark USA | Mansfield, MA | Aug 2021–May 2022

Recruited to modernize analytics across a \$1B+ enterprise with 11 acquired divisions and heterogeneous legacy ERP systems.

- Partnered with Marketing, Supply Chain, Finance, and executives to prioritize 50+ use cases and KPIs, then secured support and funding for an AWS modernization roadmap.
- Upskilled a four-person analytics team into a seven-member AWS engineering and analytics function, establishing Agile execution, risk management, reporting, and accountability.
- Delivered the company's first enterprise data lake within seven months, automating nightly ingest, cataloging, consolidation, normalization, and querying across 11 ERPs.
- Reduced cross-ERP reporting from roughly two months of manual effort to one day—a 60x improvement in enterprise insight delivery.

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

Provided architecture and engineering leadership across five enterprise Data Privacy and AI Governance initiatives, partnering with Product, Architecture, SRE, and security.

- Established cloud, security, CI/CD, and governance patterns spanning model ownership, transparency, bias and drift monitoring, auditability, human oversight, and responsible AI.
- Delivered a cloud-native privacy-rights platform within eight months, addressing CCPA, CPRA, and GDPR while automating roughly 90% of a manual workflow.
- Advanced serverless and managed-service adoption using AWS Control Tower, EventBridge, Fargate, Kafka, API Gateway, Aurora, IAM, and CDK infrastructure as code.

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

Recruited out of AWS by the NorthBay Solutions CEO to head up a new AI / ML practice and create reusable enterprise data and machine-learning offerings.

- Built the firm's AI/ML capability and advised enterprise customers on strategy, modernization, solution architecture, and production implementation.
- Won CEO approval for an AWS Data Lake + ML Accelerator, built the first working system hands-on, and then led a 24-person remote team toward commercial delivery.
- Delivered AI, analytics, and knowledge-graph solutions for customers including Dun & Bradstreet, Harvard Business Review, and Invesco.

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

Served as AWS's first Senior Specialist Solutions Architect for AI/ML after leading Big Data and Analytics, guiding enterprise adoption and technical field readiness.

- Advised 200+ enterprise customers on AI strategy, machine learning, data lakes, cloud analytics, application architecture, and production adoption.
- Founded and led global AI/ML and Big Data communities of 80 and 100 specialists, setting standards, practices, learning priorities, and delivery readiness.
- Trained 180+ solution architects, created 170+ enablement assets, and delivered 50+ conference presentations translating emerging technologies into production decisions.
- Earned AWS ML Elite Gold and Platinum recognition and contributed to AWS AI/ML and Big Data certification programs.

EDUCATION & RECOGNITION

- Ph.D., Molecular Biology / Microbiology & Immunology, University of Western Ontario; Canadian Medical Research Council Post-Doctoral Research Fellowship.
- U.S. Patents 8,793,073 and 8,489,334; BioWizard intelligent-agent patent application.

TECHNOLOGIES

AI-assisted engineering • AI agents • Agentic AI • Generative AI • LLMs • RAG • LangChain • LangGraph • AI orchestration • AI evaluation • Observability • Responsible AI • AI governance • MLOps • Model governance • AWS • Amazon SageMaker • AWS Bedrock • TensorFlow • Python • FastAPI • PostgreSQL • PGVector • SQL • Spark • EMR • Glue • Athena • S3 • Lambda • Step Functions • EventBridge • API Gateway • Kafka • Kubernetes • Docker • REST APIs • CI/CD • Data lakes • Distributed systems • Workflow automation