

RYAN CARDEN

Atlanta, GA | 470-699-4848 | ryan@RyanCarden.com

linkedin.com/in/ryanmcarden | RyanCarden.com | github.com/ryanmcarden | Remote / Hybrid

AI ENGINEERING MANAGER & HANDS-ON AI SYSTEMS ARCHITECT

PROFESSIONAL SUMMARY

AI Engineering Manager and hands-on AI systems architect with 10+ years building production software, data, automation, analytics, and business systems, with recent focus on LLM applications, RAG, AI agents, vector search, computer vision, and workflow automation. Strongest where AI has to connect to real business systems — legacy SQL databases, internal tools, customer support workflows, e-commerce/search, marketing automation, and manufacturing operations. Built AI systems from scratch inside a live manufacturing environment: RAG chatbots, SQL-connected agents, semantic site search, MCP-style tooling, computer vision pipelines, and n8n automation. Equal parts architect and leader — manages cross-functional teams of up to 10 and budgets exceeding \$2M while still writing the architecture.

CORE COMPETENCIES

AI Systems: RAG architecture, LLM applications, AI agents, vector search (Pinecone, Weaviate), prompt engineering, MCP Protocol / MCP-style tooling, chatbot architecture, structured outputs, model evaluation, computer vision, PyTorch, ViT-based pipelines

Automation & Data: n8n, PostgreSQL, Microsoft SQL Server, Pinecone, Ollama, nomic-embed-text, hybrid vector/SQL retrieval, ETL pipelines, Power BI, SQL reporting

Cloud & Infrastructure: Azure (AKS, OpenAI, Functions), AWS (SageMaker, Lambda, API Gateway), Docker, Cloudflare Access, Azure AD, Hyper-V, Windows/Linux server administration, HTTPS APIs, hybrid on-prem/cloud systems

Software & Integration: .NET, React, REST APIs, GraphQL, legacy modernization, ERP/WMS integrations, internal business tools, production support systems

Leadership: Technical roadmapping, AI strategy, vendor selection, executive stakeholder communication, \$2M+ budget ownership, cross-functional delivery, OKR alignment, Agile/Scrum

SELECTED AI SYSTEMS

Live AI Lead-Qualification Chatbot | RyanCarden.com — Built a live chatbot from scratch using n8n, OpenAI (GPT-4o), Postgres-backed conversation memory, structured JSON output, lead scoring, transcript logging, and SMS notification. Qualifies hiring/project intent, captures contact info, preserves context across sessions, and runs as a live demonstration of production AI workflow architecture.

Production RAG Support System | Stitch America — Designed and deployed a self-hosted RAG chatbot (.NET adapter + Pinecone vector store) connected to a legacy SQL CRM. Deflects approximately 50% of inbound support tickets and eliminated the need for a multi-year ERP re-platform effort. Deployed across live chat, email, SMS, and voice.

AI Hybrid Site Search | Stitch America — Built semantic site search for StitchAmerica.com using Pinecone vector embeddings, Ollama (nomic-embed-text), and SQL-based hybrid retrieval — self-hosted via n8n, near-zero inference cost, covering pages, products, and blog content.

SQL-Connected AI Tooling (MCP-style server) | Stitch America — Built a production data-access server exposing SQL Server and Postgres to AI tools over HTTPS — consumed by Claude Desktop, n8n, and a custom React backend. Secured with Cloudflare Access and Azure AD; read-only enforced at the server layer.

PROFESSIONAL EXPERIENCE

AI Engineering Manager | Stitch America, Inc. | Bremen, GA | 2017 – Present

AI Architecture & Applied AI Systems

- Designed and deployed a self-hosted RAG support system (.NET adapter + Pinecone vector store + internal business data), deflecting ~50% of inbound support tickets and eliminating the need for a multi-year ERP re-platform. Deployed across chat, email, SMS, and voice.
- Built the company's AI application stack from scratch: chatbot workflows, semantic search, SQL-connected internal tools, AI artwork review, vector retrieval, and automation pipelines integrated with live operational systems.

- Built a production MCP-style data-access server exposing SQL Server and Postgres to AI tools over HTTPS — integrated with Claude Desktop, n8n, and a custom React backend; secured with Cloudflare Access + Azure AD, read-only enforced at the server layer.
- Built an AI hybrid site search system using Pinecone, Ollama (nomic-embed-text), SQL-based retrieval, and n8n orchestration — searches products, pages, and blog content semantically on self-hosted infrastructure.
- Built a computer vision pipeline using PyTorch and ViT-based models trained on 40K+ matched embroidery design files, improving automated digitizing classification over earlier baselines (meaningful gains v1→v3).
- Deployed a Claude API-powered embroidery suitability agent to assist customer artwork intake and reduce manual review time.
- Built a multi-step AI advertising automation pipeline (n8n + GPT image generation + Gemini video analysis + GPT copy generation + Meta Graph API) — turns a product input into multiple ad creative variants and live Meta ad campaigns with minimal manual production work.

Cloud, Data & Legacy Modernization

- Modernized a 300K+ line legacy .NET application into containerized services on Azure, reducing page-load times by 48% and enabling new AI-integrated services.
- Maintained and extended 21+ years of SQL Server production data across 21M+ records, supporting executive reporting, AI agent integrations, ML feature pipelines, customer support tools, and operational analytics.
- Built internal reporting, search, and automation systems connecting SQL Server, Postgres, Pinecone, n8n, cloud APIs, and legacy business applications.

Technical Leadership & Program Ownership

- Owned technical roadmap and cross-functional delivery for AI, automation, data, infrastructure, and platform modernization initiatives.
- Managed a \$2M+ annual technology budget covering cloud services, software platforms, infrastructure, advertising technology, vendors, and internal systems.
- Led vendor selection, architecture reviews, executive workshops, cost modeling, and build-vs-buy analysis for AI tooling, cloud systems, and operational software.
- Coordinated work across engineering, marketing, operations, analytics, vendors, and executive stakeholders.

AI & Digital Consultant (Select Projects) | Freelance | Remote | Periodic engagements

- Periodic freelance engagements alongside full-time role — AI-assisted websites for professional-services clients (chatbot integration, SEO automation, lead capture), and digital infrastructure for political campaigns (web, advertising, email automation, voter targeting).

Director of Marketing & Digital Projects | Stitch America, Inc. | Bremen, GA | 2013 – 2017

- Re-architected a legacy e-commerce platform (IIS → Nginx, SQL → PostgreSQL, GA integration) and rebuilt the BI pipeline — increased organic traffic 35% and checkout conversion rate 18% year over year.
- Built SQL and Python ROAS attribution models to optimize a \$600K annual advertising budget across Google, Meta, and Microsoft Ads, increasing repeat-purchase rate 22%.
- Directed a cross-functional team of 7 across engineering, design, copy, analytics, and marketing to launch 2 new product lines, contributing \$1.2M in Q4 incremental revenue.

Genius Administrator – Apple Maps Rollout Team | Apple Inc. | Knoxville, TN | 2011 – 2013

- Triaged and resolved 200+ geospatial data defects per week during the Apple Maps global launch, collaborating with GIS engineers to improve map accuracy across 3 metro regions.
- Onboarded and mentored 12 technicians; team achieved the district's highest Net Promoter Score of 93 for 2 consecutive quarters.

EDUCATION

B.S. Business Administration – Marketing & Logistics

University of Tennessee, Knoxville | 2006