

Aric Camarata

Principal AI Engineer

Fairfax, VA (Remote) | aric.camarata@gmail.com | linkedin.com/in/ariccamarata | github.com/acamarata

PROFESSIONAL SUMMARY

Principal AI Engineer with 20+ years building distributed systems, AI infrastructure, and full-stack applications across defense, enterprise, and product companies. Deep expertise in multi-model AI routing, hybrid retrieval (RAG), federated GraphQL, and high-throughput Go backends. Founder of nSelf (self-hosted PaaS) and nClaw (multi-model AI router). Former Lead Developer at Northrop Grumman, TS-cleared (expired). Remote-only.

Past engagements span Apple, AT&T, U.S. Treasury (GSOC), Credit Karma, Walmart Labs, and Cigna — covering consumer platforms, government security operations, and enterprise health tech. Privacy-first AI is a core design principle: local inference on Apple Silicon, classifier-driven model routing, and on-premise RAG pipelines that keep sensitive data off third-party clouds. The open-source nSelf CLI deploys a full Postgres + Hasura + auth stack on a /mo Hetzner server — matching managed-DB capabilities at a fraction of the cost. Targeting senior IC, staff, and principal roles.

CORE COMPETENCIES

AI/ML, LLM, RAG, Retrieval-Augmented Generation, Embeddings, Vector Database, pgvector, tsvector, Reciprocal Rank Fusion, Cross-Encoder Reranking, Multi-Model Routing, Agent Orchestration, Claude API, OpenAI API, MLX-LM, Qwen3, Hugging Face, MCP, Prompt Engineering, Fine-tuning, Go, TypeScript, Rust, Python, PHP, RAG pipelines, BGE-M3 reranking, RRF fusion, Ollama, Gemini API, PostgreSQL, Redis, Elasticsearch, Hadoop, HBase, Cassandra, GraphQL, Hasura Federation, REST, gRPC, Docker, Nginx, Hetzner VPS, Vercel, Cloudflare, Linux, Next.js (App Router), React, Tailwind CSS

PROFESSIONAL EXPERIENCE

Principal AI Engineer — Unity

Jan 2024 — Present

Fairfax, VA (Remote)

- Architected nClaw — a privacy-first multi-model AI platform with classifier-driven tier-ladder routing across Claude, GPT-5.4, and local Qwen3 / MLX-LM on Apple Silicon, achieving ~60% inference cost reduction versus single-frontier-model baseline.
- Designed Memory Augmented Retrieval (MAR): PostgreSQL tsvector full-text + pgvector ANN + RRF score fusion + BGE-M3 cross-encoder reranking — sub-100ms semantic search with year-old conversation recall via ltree topic taxonomy and summary pyramids (turn / day / week / month).
- Built privacy-first inference layer: local models handle PII-bearing and sensitive operations; frontier models gated by a PII redaction and secret detection pipeline — full data sovereignty for regulated workloads.
- Open-sourced nSelf CLI (Go, 5,000+ lines, 30+ commands): single-command self-hosted backend stack (PostgreSQL, Hasura, Auth, MinIO, Redis, Nginx) replacing \$40–100/mo managed databases with \$4/mo Hetzner deployments.

Principal Software Architect — Unity

Jan 2014 — Jan 2024

Remote

- Founded and led Unity Platforms for a decade, delivering principal-level software architecture to Credit Karma, Walmart Labs, Cigna, Cambium Networks, Progress Software, and CMD.
- Architected Unity Sites — a multi-tenant Next.js (App Router) platform serving Flock, Veterans, FlockTV, and FlockNAS brands from a single codebase; hostname-based middleware with per-tenant feature flags, Hasura RLS, and JWT custom claims enforcing zero cross-tenant data leakage at the schema level.
- Designed distributed event-driven systems on AWS Lambda, DynamoDB, GraphQL Federation, Kafka, and SNS/SQS for financial services and healthcare clients.
- Built nSelf from an internal 200-line deployment script into a full Go CLI deployed across enterprise clients — groundwork for the AI-native platform built from 2024 onward.

Senior Developer — United States Institute of Peace

Apr 2013 — Dec 2013

Washington, DC

- Developed and hardened web applications in C# / .NET / ASP.NET for a federal peace-building institution.
- Led application security hardening — defense-in-depth access controls, credential-stuffing mitigation, and secure session management aligned to federal security standards.
- Delivered full SDLC on compressed government timelines in a zero-tolerance security environment.

Senior Developer — Double Prime

Feb 2013 — Apr 2013

Remote

- Built API-first, headless content architecture using decoupled Drupal — an early implementation of platform patterns that have since become industry standard for high-scale editorial systems.
- Integrated Akamai CDN and Alfresco content management into high-traffic global editorial platforms; developed Perl-based data transformation pipelines for large-volume content migration.

Lead Developer / Architect — Northrop Grumman

Jul 2011 — Feb 2013

McLean, VA

- Led development on the U.S. Treasury GSOC — national-scale cybersecurity observability platform for the Department of the Treasury. Held active TS clearance (since expired).
- Architected real-time big data pipeline: Hadoop, Hive, HBase, and Cassandra under 24/7 multi-site disaster recovery SLA — large-scale data engineering infrastructure that forms the conceptual foundation of modern ML pipelines.
- Led a team of 6+ engineers through full SDLC; introduced DevOps practices (Git, TDD, CI/CD environments) and reduced deployment cycles at national security scale.

Senior Developer — College Board

Apr 2011 — Jul 2011

Reston, VA

- Delivered front-end and application development on the Big Future college search platform (tier-1 national traffic).
- Built feature components on Spring MVC + GWT stack with JSP/JSTL templating.

Senior Developer — Able Engineering

Jan 2011 — Apr 2011

Phoenix, AZ

- Built employee evaluation platform on Spring MVC — form workflows, scoring logic, and admin reporting.
- Delivered full feature set from requirements to deployment on compressed timeline.

Senior Developer — Capgemini

Oct 2010 — Jan 2011

Phoenix, AZ

- Built internal IT management tools on Zend Framework: ticketing system, helpdesk portal, system monitoring dashboard, and on-call rotation tool.
- Worked alongside enterprise consulting team serving large telecom and insurance clients.

Senior Developer — Pearson Education

Jul 2010 — Oct 2010

Chandler, AZ

- Delivered CMS integrations (Drupal, WordPress, Joomla, Mura) for educational publishing platforms.
-

Built Genesys/Avaya call-center tooling for customer support operations.

Software Developer — Apple Inc.

Oct 2007 — Jul 2010

Chandler, AZ

- Built and maintained internal enterprise portal and knowledge base systems on Drupal and the LAMP stack, serving engineering and support teams across Apple divisions.
- Developed tools used daily by thousands of internal users — built to the reliability and maintainability standards of Apple's internal engineering culture.

Software Developer — AT&T

Apr 2006 — Oct 2007

Phoenix, AZ

- Led consolidation of call center support tooling as part of the AT&T / SBC AFSS merger initiative — unified NOC, Engineering, and Customer Support teams onto a single AJAX-based platform.
- Early experience with large-scale systems integration, observability tooling, and reliability engineering in a 24/7 enterprise telecom environment.

Manager — Roxbury Group

Feb 2004 — Mar 2006

Fairfax, VA

- Combined branch management and software development for a regional IT services firm — technical leadership alongside operational responsibility from the start of my career.
- Deployed and administered Windows Server networks for SMB clients; built custom inventory tracking and rental management software (C, TCL, PHP) that replaced manual processes and reduced reconciliation time significantly.

19D Cavalry Scout — U.S. Army

Mar 2001 — Feb 2004

Virginia

- Served as 19D Combat Cavalry Scout, U.S. Army.
- Service period included September 11, 2001 and the early post-9/11 operational period.
- Foundational experience in mission-critical decision-making, operational discipline, and leading under pressure.

EDUCATION

Bachelor of Science, Management Information Systems — Belltown University

Jan 2008 — Dec 2011

PA

Full Curriculum — 10 Certifications

freeCodeCamp

Completed the full freeCodeCamp curriculum: Responsive Web Design, JavaScript Algorithms & Data Structures, Front End Libraries, Data Visualization, Back End APIs, Quality Assurance, Python for Scientific Computing, Python for Data Analysis, Information Security, and Machine Learning with Python.

Computer Science & Mathematics — Full Library

MIT OpenCourseWare

Self-directed study through MIT OpenCourseWare covering algorithms, data structures, mathematics, and computer science fundamentals.