

CAM AN NGUYEN THANH

Senior Fullstack Engineer · Backend, Frontend & AI/LLM

caman.nguyenthanh@gmail.com • github.com/CamAnNguyen • [LinkedIn](#) • [Google Scholar](#)

PROFILE

Senior fullstack engineer with 15+ years building and shipping SaaS platforms and data-intensive products end to end — strong across backend (Ruby on Rails, Node.js, ASP.NET, PostgreSQL/MongoDB) and frontend (React, Redux). For the last few years I've built hands-on with AI and LLMs: retrieval-augmented generation (RAG) over business data, AI customer-support agents, model routing by cost, and local inference with Ollama and llama.cpp. Earlier, I scaled an affiliate-tracking platform to 100M+ daily visitors — cutting its Shopify app p95 LCP from 20s to 2.1s (earning the Built for Shopify designation), making data APIs ~7× faster, and lowering AWS costs by 20%. I take products from prototype to production, communicate clearly with both technical and non-technical stakeholders, and move fast in ambiguous environments. Author of published open-source Ruby libraries, including native C++ bindings for the RDKit and Open Babel cheminformatics toolkits.

CORE SKILLS

Backend & Systems: Ruby on Rails, Node.js, Express.js, ASP.NET / C#, Python, PHP, Sidekiq, REST APIs

AI & LLM Engineering: OpenAI & Claude APIs, RAG & vector search, AI agents, prompt/context tuning, local inference (Ollama, llama.cpp), model routing & usage tracking, AI-assisted dev (Cursor, Windsurf)

Databases & Data: PostgreSQL, MongoDB, MySQL, Redis, Firestore, query optimization, large-scale data retrieval

Cloud & DevOps: AWS, Azure, Docker, Kubernetes, GitLab CI/CD, Grafana, Prometheus, Azure Monitor, observability

Frontend & Product: React, Redux, WebXR, IntercoolerJS, interactive 3D/VR interfaces, product performance

Specialized Domains: GDPR-compliant tracking, document & information extraction, cheminformatics, molecular search, high-throughput packet inspection

Languages: Vietnamese (native), English (professional working proficiency)

PROFESSIONAL EXPERIENCE

Senior Software Engineer · Contract

Remote · 2020 – Present

Affiliate Tracking SaaS — E-commerce Integrations, Performance & Scale

- Built and scaled an affiliate tracking SaaS platform serving 100M+ daily visitors across Shopify, WooCommerce, BigCommerce, ReCharge, and related e-commerce integrations.
- Reduced 95th percentile Largest Contentful Paint (LCP) from 20s to 2.1s, helping secure the Built for Shopify designation and materially improving merchant experience.
- Redesigned backend architecture and optimized complex MongoDB queries, making large-scale data-retrieval API responses roughly 7× faster.
- Rebuilt application observability and monitoring to expose bottlenecks, support rapid incident response, and sustain high performance after release.
- Partnered with product and operations teams to redesign retention and storage, cutting database size by 50% to a 2 TB rolling 18-month active window in RDS and lowering AWS infrastructure costs by 20%.
- Refactored a large legacy Rails codebase and increased automated test coverage by 50%, improving reliability, maintainability, and release confidence.
- Designed GDPR-compliant tracking mechanisms and maintained Grafana-based monitoring and logging pipelines.

Tech: Ruby on Rails, Sidekiq, MongoDB, MySQL, Redis, AWS, Grafana, IntercoolerJS

AI Customer-Support & CRM Platform — RAG, Model Routing & Pricing

- Designed a retrieval-augmented generation (RAG) pipeline that turns CRM records into a live knowledge base, grounding LLM answers in each client's own products, pricing, and policies.
- Built an AI customer-support assistant for Facebook Page businesses that automates first-line buyer responses and cuts manual reply load for the merchant's team.
- Implemented LLM model routing that maps each customer to a model tier by price and performance, with per-customer AI usage tracking for cost control and billing.
- Designed and implemented the underlying configurable CRM pricing platform supporting complex discount and sales-strategy requirements.
- Architected CI/CD and monitoring from scratch (Grafana, Prometheus, Azure Monitor) and owned mobile delivery end to end, including resolution of complex Apple App Store review issues.
- Hired to build the engineering team from zero — recruited and onboarded 6 engineers, set delivery workflows and engineering standards, then served as ongoing technical advisor.

Tech: C#, ASP.NET, Azure, OpenAI/LLM APIs, RAG, vector search, Grafana, Prometheus, CI/CD

VoicePing WebXR Virtual Tour & Workspace — Real-Time Multilingual VR

- Shipped immersive WebXR features for virtual tours, distributed collaboration, and travel exploration at startup speed during the company's stealth phase — work that helped VoicePing reach a launched, surviving product.
- Integrated VoicePing's live-translation API into multilingual VR tour workflows, letting participants follow tour guides in their preferred language in real time.
- Built application-layer workflows for multilingual tour sessions, including UI interactions, client synchronization, real-time state handling, and backend support for interactive 3D/VR experiences.
- Integrated Mapbox-styled 3D maps with OpenLayers offscreen-canvas rendering to move heavy processing off the main thread and hold frame rates steady in AR/VR.
- Designed backend services for heavy 3D assets, session state, and distributed client coordination.

Tech: Node.js, Express.js, React, MySQL, Firebase/Firestore, Redis, AWS, WebXR, VoicePing API

NativeX (Ed-Tech) — App Support & Speech-AI Integration During Acquisition

- Joined the team for a six-month engagement to maintain, troubleshoot, and support the language-learning application through the company's transition and acquisition.
- Maintained and improved the integration with LanguageConfidence's speech-AI API, which powers automated pronunciation and spoken-spelling evaluation for learners.
- Kept the app stable and shipping during organizational change, coordinating fixes and support directly with the in-house and acquiring teams.

Tech: LanguageConfidence Speech-AI API, REST APIs

Research Collaborator (Part-time) — ComPlat Group, KIT

Remote · 2016 – Present

ChemScanner — Open-Source Research Tool for Chemical Information Extraction

- Collaborate part-time with the ComPlat group at KIT (where I completed my PhD) on ChemScanner, an open-source tool that extracts and interprets chemistry reactions from scientific documents — an ongoing research collaboration that has produced peer-reviewed publications (J. Cheminform., 2019).
- Extended the pipeline with locally-hosted LLM inference (Ollama, llama.cpp), tuning prompts to interpret and normalize chemical information the deterministic extractor cannot resolve — kept fully on-premise to protect research data.
- Built and maintain the document-management and validation interfaces, integrating Ruby/Rails and React with C/C++ extraction components via SWIG.

Tech: Ruby on Rails, React, C/C++, SWIG, Ollama, llama.cpp, local LLM

Research Engineer — Karlsruhe Institute of Technology

Germany · 2016 – 2020

Chemotion ELN — Electronic Lab Notebook for Chemists

- Contributed to core full-stack development of Chemotion ELN — a production electronic lab notebook in active use since 2015 for managing private experimental data at lab scale across research institutes and universities.

- Engineered specialized molecular search capabilities for exact, structural, and substructural queries over large chemical datasets.
- Automated computational chemistry workflows with intuitive web interfaces and built GitLab/Docker CI/CD pipelines.
- Applied ML-based methods to generate 3D molecular coordinates from 2D structures, automating a step that previously required manual setup.

Tech: Ruby on Rails, React, Python, Docker, GitLab CI/CD

Research Associate — Infnis Inc.

Seoul, South Korea · 2014 – 2016

PIUS — High-Speed Packet Inspection in User Space

- Migrated the Linux kernel TCP/IP stack into user space, enabling sustained raw packet inspection close to 10 Gbps for dangerous-pattern detection without compromising application stability.

Tech: C/C++, Boost

Software Engineer — Cargigi (acquired by eBay Motors)

Ho Chi Minh City · 2012 – 2014

ADAM — Auto Data Administration and Marketing

- Built customer-facing inventory management interfaces for auto dealer networks.
- Integrated automated web services to publish high-volume dealership advertisements to Craigslist.
- Developed customized internal CRM tools to track operational metrics and improve workflow visibility.

Tech: ASP.NET

System Engineer — VNG Corporation

Ho Chi Minh City · 2011 – 2012

World of JX — MMORPG Infrastructure

- Administered and scaled production infrastructure for one of Vietnam's most popular MMORPGs, maintaining availability and performance across 200+ dedicated servers.
- Automated log parsing, maintenance tasks, and internal request workflows through custom monitoring and administration interfaces.

Tech: PHP/Codelgniter, Zabbix, Shell scripting

EDUCATION

Ph.D., Organic Chemistry — Karlsruhe Institute of Technology, Germany	2016 – 2019
M.Sc., Aerospace Information Engineering — Konkuk University, South Korea	2014 – 2016
B.Eng., Computer Science and Engineering — Ho Chi Minh City University of Technology, Vietnam	2006 – 2011

SELECTED PUBLICATIONS

-
- **Nguyen, A.**; Huang, Y.-C.; Tremouilhac, P.; Jung, N.; Bräse, S. ChemScanner: Extraction and re-use(ability) of chemical information from common scientific documents containing ChemDraw files. J. Cheminform. 2019, 11, 77.
 - Huang, Y.-C., Tremouilhac, P., **Nguyen, A.** et al. ChemSpectra: a web-based spectra editor for analytical data. J. Cheminform. 2021, 13, 8.
 - Tremouilhac, P.; **Nguyen, A.**; Huang, Y.-C.; Kotov, S.; Lütjohann, D. S.; Hübsch, F.; Jung, N.; Bräse, S. Chemotion ELN: An open-source electronic lab notebook for chemists in academia. J. Cheminform. 2017, 9(1), 54.

Full publication list available on request.

EXTERNAL LINKS

[Google Scholar](#) | [Stack Overflow](#) | [ResearchGate](#) | [GitHub](#)