

Amer A Qresh

Full-Stack & AI Engineer

qreshamer@gmail.com | Village Dr, Edison, NJ 08817, USA | <https://github.com/kame-hash> | www.amerqresh.com/portfolio

PROFESSIONAL SUMMARY

A cross-cutting specialization layered on top of Python, Rails and Node backends: LLM integrations, RAG pipelines, multi-agent orchestration, vector/semantic search and AI-driven analytics, shipped into real production products.

TECHNICAL SKILLS

Core Backend: Ruby on Rails 6/7, Ruby 3.x, ActiveRecord, ActionCable, ActionMailer, Sidekiq, Resque, Devise, Pundit

Frontend: JavaScript (ES6+), TypeScript, React.js, Next.js, Vue.js, Nuxt.js, Svelte, PWAs, Tailwind CSS, HTML5 / CSS3

ML / AI: TensorFlow, Python ML pipelines, LangChain, LlamaIndex, OpenAI GPT-4, RAG pipelines

LLMs / APIs: OpenAI API (GPT-4), OpenAI Whisper, ChatGPT integration

CI/CD: GitHub Actions, CircleCI, Capistrano, Vercel, Netlify, Heroku

Databases: PostgreSQL, MySQL, MongoDB, Redis, Schema design & query optimization

Cloud / DevOps: AWS (EC2, S3, RDS, Lambda, CloudFront, VPC), Docker, Kubernetes, GitHub Actions, CircleCI, Heroku, Capistrano

Vector / Search: pgvector semantic search, Vector search, PostGIS-driven recommendations

Testing: RSpec, FactoryBot, Capybara, SimpleCov, VCR, Minitest, TDD / BDD (90%+ coverage)

Data / BI: Snowflake, Streamlit, ETL Pipelines, pgvector, PostGIS, High-Volume Optimization, Predictive Analytics

Web3: Web3, DApps, Blockchain, Smart Contracts

PROFESSIONAL EXPERIENCE

Tkxel — Senior Full Stack Software Engineer

Feb 2024 - Continue

- Owned end-to-end delivery of RAG-powered search and knowledge platform (MicroHealth) from problem definition and pipeline architecture through to production, iterating on ranking, latency, and response quality based on real usage patterns
- Designed scalable document ingestion, retrieval, and LLM-based response pipelines balancing performance vs accuracy trade-offs for real-world healthcare data usability
- Integrated ML models for predictive analytics and AI-driven decision-making into existing production systems, owning the full pipeline from model integration to live inference
- Worked directly with stakeholders to refine AI output relevance and align LLM responses with real user needs
- Built AI-powered multi-tenant SaaS marketplace (GoZupees) with automation workflows, service orchestration, and intelligent decisioning layers
- Data Engineering:
 - Architected a production-grade ETL data pipeline using Python, AWS S3, and Snowflake handling large-volume data feeds with fault-tolerant design
 - Designed data ingestion and enrichment workflows supporting downstream ML model consumption and analytics
- Backend Development:
 - Built high-performance Python (FastAPI/Django) and Rails backends powering AI-driven SaaS platforms with clean REST API design
 - Designed multi-tenant architectures enabling flexible and scalable AI feature expansion across products

- Integrated Plaid, Bloom, Stripe, and Tilled for complete financial and payment user journeys within AI-powered platforms
- DevOps & Infrastructure:
- Managed Dockerized deployments with Kubernetes orchestration and GitHub Actions CI/CD pipelines supporting continuous AI model delivery
- Provisioned AWS infrastructure (EC2, S3, RDS, Lambda) optimized for ML workloads and high-volume data processing
- Implemented Sentry monitoring and automated testing frameworks ensuring AI pipeline reliability in production

Strategic Systems International — Senior Software Engineer

November 2021 – January 2024

- Built scalable Node.js (plus Rails) apps with React/Vue.js frontends.
- Designed microservices with modular, independently deployable services and secure APIs.
- Built CI/CD pipelines, ETL (AWS S3 + Snowflake) and LLM/AI integration.
- FrostyMetrics is an analytics platform built for Avalanche, providing developers with APIs and SDKs to access blockchain data. It supports real-time metrics, historical data, and custom analytics for applications on Avalanche.
- Designed and implemented CI/CD pipelines, streamlining deployment workflows and release processes.
- Built ETL data pipeline integrating from Python project to Snowflake for high-volume data processing via AWS S3.

Solo — Senior Devops Engineer

Sep 2019 - Oct 2021

- Cloud Infrastructure & Architecture
- Architected cloud-native deployments on AWS and GCP with high availability, fault tolerance, and disaster recovery configurations
- Provisioned and managed AWS core services (EC2, S3, RDS, Lambda, VPC, CloudFormation) using Infrastructure as Code principles
- Designed serverless and microservices architectures optimized for scalability, cost-efficiency, and independent deployability
- CI/CD & Containerization
- Built and maintained end-to-end CI/CD pipelines via GitHub Actions across dev, staging, and production environments
- Managed Dockerized deployments with Kubernetes orchestration ensuring zero-downtime releases and auto-scaling
- Built Playwright browser automation pipelines and automated testing frameworks improving reliability across CI pipelines

The Tie — Software Engineer

July 2018 – August 2019

- Built a financial data and analytics platform delivering real-time market insights and AI-driven analytics
- Developed secure authentication and user management with Auth0 integration
- Designed scalable backend architecture using Ruby on Rails
- Built responsive frontend using Vue.js
- Deployed and managed infrastructure using Docker and AWS

RYVYL - Quick Card — Software Engineer

March 2017 – June 2018

- Shipped full-stack apps with React, Vue.js and Nuxt.js (plus Rails) REST APIs and third-party services.
- Scaled a REST API for 2M+ users on a social trading platform; 99.95% uptime SLA under sustained load.
-

Ran Sidekiq + Redis background jobs at 500,000+ async tasks/day (trade notifications, portfolio recalcs, email workflows) with zero data loss.

- Led Heroku → AWS EC2/RDS migration; 30% monthly hosting cost.

PROJECTS

GoZupees

Tech: Python (FastAPI), PostgreSQL (pgvector), OpenAI GPT-4, LangChain, SQLAlchemy, Twilio, WebSockets, React, TypeScript, Docker, AWS

- Real-time AI call center management platform; WebSocket-based AI suggestions at 1–2s latency (OpenAI GPT + pgvector RAG).
- Structured agent/ticket management with full CRUD REST APIs (FastAPI + SQLAlchemy).
- Twilio WhatsApp, live call transcription and Stripe subscription management.
- Docker on AWS with CI/CD via GitHub Actions.

Genie AI

Tech: Python (FastAPI), PostgreSQL (PostGIS), LangChain, OpenAI API, Redis, RabbitMQ, Docker, Auth0, SQLAlchemy

- Scalable location-based discovery APIs with PostGIS analytics.
- Real-time chat via WebSockets; RabbitMQ event messaging.
- AI-powered recommendations (LangChain LLM + RAG personalization).
- Auth0 auth and Redis caching supporting thousands of concurrent users.

SoloSuit

Tech: React, Node.js, Playwright, Plaid, Bloom API

- Led integration of Plaid + Bloom into the platform.
- Introduced Playwright testing from scratch for critical debt-law workflows; delivered features and bug fixes for platform stability.

The Tie

Tech: ChatGPT, Ruby on Rails, Vue.js

- Real-time market insights with ChatGPT integration.

TribeAWL

Tech: React.js, Tailwind, Ruby on Rails, PostgreSQL, WebSockets, Stripe, Tilled, Active Storage

- Multi-tenant SaaS: client/staff management, appointment booking, inventory tracking, role-based portals and real-time notifications.
- Stripe + Tilled PCI-compliant payments; Twilio SMS and in-app chat via WebSockets.

Frosty Metrics

Tech: Node.js, Python (Django, ML), React, Next.js, PostgreSQL, Redis, Docker, Kubernetes, AWS, WebSockets, GraphQL, LLM

- High-performance real-time crypto analytics: token tracking, historical data viz and AI-driven predictive analytics via Python ML pipelines.
- Scalable microservices, real-time data pipelines and containerized cloud deploys handling large transaction volumes; secure GraphQL APIs.

MicroHealth

Tech: Python (FastAPI/Django), PostgreSQL (pgvector), OpenAI GPT-4, LangChain, LlamaIndex, Azure, Vue.js, Streamlit

- End-to-end RAG-powered intelligent document search for healthcare data.
- Scalable Django/FastAPI ingestion pipelines with pgvector semantic search.
- Streamlit dashboards for admin analytics and model-performance monitoring.
- Iterated on relevance, latency and quality against real clinical user needs.

EDUCATION

Master's Degree

CERTIFICATIONS

- AWS Certified Solutions Architect Associate
- Certified Web Developer