Amine Belkadi

Former Founder & CTO | Engineering Manager | Staff Fullstack Engineer | Staff Devops Engineer | Senior Fullstack Engineer

Email: ajbelkadi@gmail.com | **Phone:** (929) 469-9353
LinkedIn: linkedin.com/in/belkapi **| Jobtie:** jobtie.ai/cv/amine-belkadi
Location: New York, NY (Hybrid)

Professional Summary

Engineering Leader and Software engineer with over 15 years of experience in React/Next.js, Python and Node.js ecosystems, specializing in the design of scalable backend services (REST/gRPC) and high-performance data systems. Advanced experience in leveraging AI (LLMs, RAG) for production-ready solutions and automating complex engineering workflows. Proven track record in leading cross-functional teams, implementing robust DevOps processes (IaC, CI/CD), and maintaining high engineering standards to deliver mission-critical software in Fintech and SaaS environments.

Technical Skills

- **Languages:** Python, TypeScript, JavaScript (Node.js), PHP, SQL, Java.
- **Data Engineering:** Amazon Redshift, AWS Glue, AWS Athena, ETL Pipelines, Big Data Ingestion, Pandas, NumPy, Airflow, Kafka.
- **Frameworks & APIs:** FastAPI, Flask, Django, Next.js, React, Laravel, WebSockets, gRPC, Protocol Buffers, Prisma, Knexjs
- **Machine Learning:** LLMs (Custom Training & RAG), PyTorch, TensorFlow, Reinforcement Learning (PPO), XGBoost, LangChain, Vector DBs.
- **Infrastructure:** AWS (EC2, RDS, S3), GCP, Terraform (IaC), Docker, Kubernetes, RabbitMQ, Redis, Linux.
- **Databases:** PostgreSQL, MySQL, MariaDB, MongoDB, Redis.

Professional Experience

April / Engineering Manager / Fulltime
Series B
[70% coding hands on / 30% managerial]
September 2025 – Present | New York, NY

- **Built and led a high-performing senior full-stack engineering team**, overseeing the end-to-end recruitment process and conducting technical interviews to secure top-tier talent.

- **Architected comprehensive tax software solutions** for tax professionals and DIY filers, utilizing **TypeScript**, **Python**, and **gRPC** with **Protocol Buffers (Protos)** to ensure high-performance service communication.
- **Developed a specialized "Tax OS" support tool** for secure data administration and PII (Personally Identifiable Information) protection, implementing **Privileged Access Management (PAM)** protocols.
- **Engineered LLM-powered systems** to interpret complex IRS tax forms and automatically convert unstructured data into structured formats, leveraging **BigQuery** and **Spanner DB** on **Google Cloud Platform (GCP)**.
- **Managed cloud-native infrastructure** deployments using **Terraform**, ensuring scalable and resilient environments for data-rich, high-impact products.
- **Provided technical mentorship** to senior and staff-level engineers, elevating team capabilities and fostering a culture of architectural excellence across distributed systems.

  Stacks/Libraries/Frameworks:
  GCP, Python, ReactJS (Typescript), SpannerDB, BigQuery, turborepo,Redis, Pub/Sub, gRPC, Datadog, Langchain, Terraform

Revv: Staff Fullstack Engineer /Staff DevOps Engineer / Fulltime
Series B
[90% coding hands on / 10% team management]
January 2025 – September 2025

- I Built a fullstack ADAS dashboard for our customers with filters that automatically parses PDF car scan results and turns them into an invoice after fetching the parts needed in our DB and mapping the pricing and time to delivery.
- I improved the AWS setup by adding auto scaling to RDS (PostgreSQL), getting rid of massdriver layer for IAC and rewriting the IAC in pure terraform.
- I improved the deploy time by creating modular testing
- I set up Datadog (in code and dashboard) with slack alerts for observability although Cloudtrail was enabled. Datadog allows for better monitoring and customizing the ingestion, and alerts delivery.
- Built Claim Builder, a feature that helps streamline every insurance claim turning repair documentation into clear claim files to reduce approval time. (nextjs/python/langchain)
- I Built a database migration framework using Knex.js and PostgreSQL with temporal data tracking and automated schema versioning.

  Stacks/Libraries/Frameworks:
  Python, FastAPI, Nextjs (Typescript), PostgreSQL, KnexJS, Redis, Datadog, Posthog, Terraform, Helm, Langchain

TradeUI: Fintech Startup, **CTO/Founding Engineer**
Seed

[75% coding hands on / 15% managerial]
January 2018 – 2022

- Spearheaded zero-to-one product design and end-to-end full-stack development of the core options analytics SaaS, building from initial concept to a fully deployed production environment (ReactJS, Bun, PostgreSQL) processing millions of live events in under three months.
- Architected a high-frequency fintech platform processing terabytes of financial data for real-time market analysis.
- Engineered a Python ingestion engine to normalize massive real-time data feeds from NASDAQ and OPRA.
- Applied ML techniques using PyTorch and XGBoost for time-series forecasting, building feature engineering and model evaluation pipelines across social sentiment, news streams, and historical market data.
- Implemented Proximal Policy Optimization (PPO) algorithms for HFT and swing trading.
- Designed scalable microservices using FastAPI to serve real-time signals via WebSockets to thousands of concurrent users.

  Stacks/Libraries/Frameworks:
  ReactJS, Nextjs (Typescript),Python, FastAPI, MariaDB/MySQL, Redis, Tensorflow, Pytorch, XGBoost, Highcharts, SQS, Websockets, Socket.io, NodeJS, ExpressJS, Terraform, Grafana

FastCompany: **Engineering Manager & Senior Software Engineer** / Fulltime
Small/Midsize
[60% coding hands on / 40% managerial]
September 2018 – January 2025

- Led development of high-performance analytics and CMS platforms using Node.js, React, Redis, and WordPress.
- Pioneered Generative AI and LLM/RAG integration for content indexing and search, significantly increasing user engagement.

  Stacks/Libraries/Frameworks:
  NextJS, Typescript, PHP, Wordpress, MariaDB
  New Relic, Aurora (Postgres), Elasticache (Redis), Elastic Container Registry (ECR), ECS (Fargate), Lambda, OpenSearch, Amazon SES

Inc Magazine: **Senior Software Engineer** / Fulltime
Small/Midsize
[100% coding hands on]
July 2014 – January 2025

- Directed a transformative migration of Inc.com to a modern stack (React, Next.js, Laravel), slashing page load times by 25%.
- Built "Kudos" (licensing e-commerce platform) from the ground up, coding the full stack and configuring secure server environments.

- I built a full-fledged analytics system (Tempo) with A/B testing and an algorithm to automatically predict and select with headline will perform the best after collecting 5 minutes worth of data. This involved setting up the infrastructure on AWS, it uses redis, SQS for events, RDS (MariaDB), lambda, real time websockets, Auth0 for user roles management.
  Built in ReactJS and ExpressJS, data crunching and archiving is done using python.

  Stacks/Libraries/Frameworks:
  ReactJS, PHP, Golang, Wordpress, MariaDB
  New Relic, Aurora (Postgres), Elasticache (Redis), Elastic Container Registry (ECR), ECS (Fargate), Lambda, OpenSearch, Amazon SES,Ansible

Mansueto Ventures owns both Inc and Fastcompany I initially worked for Inc then ended up serving Inc and FC.

SproutSMS: **Staff Fullstack Developer** / Contract
Privately Funded

- I built an enterprise-grade cloud communication platform (99.9% uptime) with a unified API layer supporting Twilio, SignalWire, and Telnyx.
- Designed a MySQL database cluster and a RabbitMQ-based distributed messaging system for high-throughput SMS processing.

Stacks/Libraries/Frameworks:
ReactJS, Python, MariaDB/MySQL, Redis, Highcharts, twilio sdk,telnyx (voip), Websockets

---

AI & Language Model Projects

April (Tax for Professionals) | Engineering Manager

- Wrangling IRS paperwork isn't for the faint of heart. I built a custom document parsing pipeline one that grabs messy, unstructured text and sorts it out with specialized language models. Thanks to an early-warning validation loop I designed, we caught compliance errors before they could sneak into production. This system checks for mismatched line items and conflicting details across giant multi-page files, saving headaches later on.
- The data was a mess and schemas all over the place, random columns, shifting layouts in PDFs. I automated the whole process, mapping these chaotic tables into clean JSON for our downstream accounting tools. And about those hallucinations everyone worries about? I reworked how we chunk documents and filter metadata, so now the system extracts data from dense government records with complete accuracy.
- **Tech Stack:** Python, LangGraph, LlamaIndex, OpenAI API, PDFPlumber, JSON Schema

JobHired (Agentic Automation Platform)

Staff Software Engineer, Fullstack & AI Agents

- Designed a multi-step agentic pipeline in Python/LangChain: job URL in, LLM field mapping + browser tool-use out. Decoupled the control plane from execution with Redis Streams EDA. API publishes apply events, headless workers consume them, emit status, and scale independently.
- Treated prior answers as retrieval context before prompting. Built evals with a regression matrix that reruns every agent version against a fixed set of past job applications and scores field mapping accuracy, catching model regressions and hallucinations in CI before they reached a real application.
- **Tech Stack:** Python, LangChain, LLM APIs (OpenAI, Groq, Claude), Evals, Redis Streams, Playwright, MariaDB, Docker, GitHub Actions, FastAPI, Next.js, Bun, TypeScript, Git Hooks, GitHub API, Vector Database, Commander.js, Agentic AI Systems, RAG, Multi-Agent Orchestration, Model Evaluation Frameworks, Vector Database

Reviewcode (Adopted within April workflows)

Staff Software Engineer, AI Agents

- I wanted code reviews to actually help, not slow us down. So I hooked into the GitHub API, pulled in huge histories of pull requests, and trained the AI to review code the way our team actually writes. To make it smarter, I spun up a vector database that links code changes directly to real reviewer comments and bug fixes from the past.
- The whole process now kicks in the second you try to commit, thanks to a CLI tool I built running in your local Git pre-commit hook. Engineers get instant

feedback right on their machines in the terminal (CLI). This solution saves 30 to 90 minutes per developer on each PR lifecycle.
- **Tech Stack:** Bun, TypeScript, Git Hooks, GitHub API, Vector Database, Commander.js, LLM APIs (Claude)

TradeUI Agents | Lead AI Developer

- Trading signals from everywhere news, technicals, momentum, social sentiment and risk management. So I built a "Chain of Thought" pipeline where multiple AI agents run in parallel, each with its own specialty. I wrote a deterministic "boss agent" that uses weighted logic to pick a side, so the user gets one clear recommendation.
- I built a custom alignment algorithm that scores how much these agents agree, from 0 to 100, so the confidence is clear before the user pulls the trigger. And I got inference time down by connecting the Alpaca API straight into a large open source model via Groq—live market data.
- **Tech Stack:** React, TypeScript, Tailwind CSS, Highcharts, Groq SDK, Alpaca Markets API

Using Agents to turn unstructured data into structured

April (Tax for Professionals)

- The platform had to process 25 different tax forms and random documents all at once. I built the orchestration using ADK (adk.dev) to manage the agent workflows and set up strict JSON form templates.
- The agents map the unstructured text straight into the schemas and run validation loops to make sure the inputs are totally valid before we even touch the database. Adk handles the complex context and routing so we never drop data during the extraction process.
- Tech Stack: Python, Adk, Agent Workflows, JSON Schema, Pydantic

Revv (Auto Industry SaaS)

- We had an absolute mess of unstructured invoices and car damage reports coming in. I used Langchain to build a custom extraction pipeline that pulls this unstructured text and forces it into clean JSON that perfectly matches our strict Zod schemas.
- Once the data passes the Zod validation it gets stored and indexed making the whole thing researchable and highly queryable. I wired up the Langchain retrievers so the team can easily search through thousands of damage reports in milliseconds.
- Tech Stack: TypeScript, Langchain, Zod, Vector Databases, LLM APIs

Forward Deployed Engineering & Contracting

I worked as the person in the room who could turn what a client wanted into something that actually shipped. That meant spending real time with non technical stakeholders, asking a lot of questions, and treating the first version of the problem as wrong until proven otherwise. I moved fast in situations that didn't have a clear spec, and I stayed close enough to the client's day to day that I could tell when I was solving the wrong thing.

Event-Driven Architecture (EDA)

Developed scalable Event-Driven Architectures (EDA) to build decoupled, high-throughput distributed systems, applying deep understanding and experience across AWS SQS, Redis (Pub/Sub and Streams), Apache Kafka, and GCP Pub/Sub.

Education & Certifications

- **AWS Certified Machine Learning - Specialty** (2025)
- **University of Central Florida**:
  Bachelor of Science in Management/Entrepreneurship (2010)
- **Computer Programming Institute, Fes, Morocco**:
  Software Development (1999)