Prateek Mulye

Senior Full-Stack Engineer · Backend-heavy · Distributed Systems · Applied AI

prateekmulye@gmail.com | +1 914-426-0562 | Santa Barbara, CA | open to international relocation linkedin.com/in/prateekmulye | github.com/prateekmulye

PROFESSIONAL SUMMARY

I have spent 11+ years building back-end systems that hold up under real load, and the last stretch bringing AI into that same engineering discipline. The backbone is distributed systems: event-driven services on Kafka, 30M+ row PostgreSQL, idempotent recovery, and the service contracts that keep things honest when traffic spikes, learned first on financial and banking platforms where a mistake costs money. More recently I've been working on AI-assisted applications at Agilent and building multi-agent and retrieval systems in my own projects, and I use agentic coding agents in my daily engineering. I am open to relocating internationally for the right team.

TECHNICAL SKILLS

- **Languages:** Java (8/17/21), Elixir/Phoenix, Python, TypeScript, SQL
- **Distributed Systems & Messaging:** Apache Kafka, RabbitMQ, IBM MQ, AWS Kinesis, Oban Pro
- **AI/ML:** LLM APIs, LangGraph, LangChain, retrieval-augmented generation (Pinecone), Pydantic schema validation, response guardrails
- **Data & Storage:** PostgreSQL, AWS DynamoDB, MongoDB, Redis, Oracle
- **Cloud & Infrastructure:** AWS (Lambda, S3, EC2, Kinesis, EMR), Azure (App Services, AKS), Kubernetes, Docker, Terraform, Argo CD, Helm
- **CI/CD & Tooling:** GitHub Actions, Jenkins, LaunchDarkly
- **Web & APIs:** REST, GraphQL (Apollo, Absinthe), React (TypeScript), Next.js
- **Observability:** OpenTelemetry, Datadog, CloudWatch, Splunk, Sentry
- **Testing:** JUnit 5, ExUnit, Mockito, Testcontainers, Cypress

PROFESSIONAL EXPERIENCE

Software Engineer, Manufacturing Systems | Agilent Technologies, Carpinteria, CA · Jan 2026 – Present

- Contributing to a manufacturing-AI effort that is replacing Excel-heavy plant workflows with grounded, AI-assisted applications.
- Prototyping a retrieval-augmented (RAG) layer over standard-operating-procedure documents, with schema validation and guardrails so the output is auditable on a factory floor, not just plausible.
- Scoping the first pilots directly with manufacturing and quality engineers, picking the workflows where automation actually moves the needle.

Senior Software Engineer | HG Insights, Santa Barbara, CA · Sep 2021 – Nov 2025

- Designed and owned a geo-standardization service that normalized 15M+ U.S. and Canada company records into canonical shapes. It became the framework every downstream ingestion and UI workflow runs through, carried through a phased rollout with the EVP of Data Solutions, and it lifted data quality and customer trust.
- Scaled PostgreSQL past 30M rows with partitioning, composite indexing, and read replicas, and cut customer-facing dashboard latency by more than 60% while holding p99s steady through reporting bursts.
- Automated the company-spend refresh across corporate hierarchies with scheduled, on-demand, and field-triggered Oban Pro jobs. A manual monthly batch became a weekly, hands-off cadence, and the analysts stopped chasing stale companies by hand.
- Built a contract-processing pipeline over 80K+ records in Elixir and Phoenix on Oban Pro, idempotent and fault-tolerant, that cut reporting turnaround by 60% without overloading shared compute.
- Set the team's default for observability with OpenTelemetry and Datadog, the tracing, SLO dashboards, and queue-depth alerts the platform leaned on for triage, and mentored the mid-level engineers behind it.

Senior Backend Engineer | Cognizant (Client: Capital One), Vienna, VA · Nov 2018 – Sep 2021

- Led the strangler-pattern decomposition of a legacy Java monolith into modular, contract-first services, which let teams ship in parallel instead of tripping over a shared codebase.
- Owned the Kafka workflows where correctness mattered most (Spring Boot, DynamoDB): a dead-letter-queue strategy, bounded retry with backoff, and idempotent replay, so operators could recover from a bad message without double-processing a financial transaction.
- Built an OpenResty (NGINX and Lua) service-virtualization layer that held 12K+ TPS in test, so integration testing stopped being the thing that gated a release.
- Built Lambda processors that ingested, transformed, and validated financial-transaction data across asynchronous flows, sized and rate-limited to keep them from starving shared resources.

Software Engineer | AurionPro Solutions, Pune, India · Mar 2013 – Sep 2015

- Built the secure login and anti-phishing layer for a Kenyan bank's first online-banking platform, in Java under real regulatory and latency constraints. It was greenfield, and it had to be right the first time.

Earlier Experience

- **Software Engineer, Tuutkia** (San Jose, CA · 2018) and **Software Engineer Intern, Illinois Department of Public Health** (Springfield, IL · 2017–2018), the latter while finishing the M.S. in Computer Science.

KEY PROJECTS

FinResearch AI (2026)

- A multi-agent market-intelligence platform that aggregates, analyzes, and synthesizes financial data on its own, with Pydantic-validated outputs and guardrails on every step. Live: huggingface.co/spaces/prateekmulye/FinResearchAI

EDUCATION

- **M.S. Computer Science**, University of Illinois at Springfield, USA · 2018
- **Post-Graduate Diploma, Advanced Computing**, C-DAC, Pune, India · 2013
- **B.Sc. Computer Application**, University of Pune, India · 2012

CERTIFICATIONS

- **Generative AI with Large Language Models**, DeepLearning.AI and AWS · Feb 2026. Verify: coursera.org/account/accomplishments/verify/NZRBB2IX3UTV

LANGUAGES

English (C2) · German (A1-A2) · Italian (A1/A2) · Hindi & Marathi (Native)