

SAHIL DEEPAK KALGUTKAR

Boston, MA | sahilkal717@gmail.com | +1 857-799-0189 | linkedin.com/in/kalgutkarsahil

PROFESSIONAL SUMMARY

- Software Engineer with 5+ years designing, building, and testing scalable, cloud-native backend services and platform tools used by internal teams and end customers.
- Experience owning features end-to-end — clarifying ambiguous requirements, designing solutions, implementing, testing, and operating in production — while partnering with Product, Compliance, and cross-functional stakeholders.
- Track record of migrating legacy systems into resilient microservices, conducting code reviews, and driving automated test coverage, coding standards, and operational best practices across teams.
- Comfortable communicating technical designs and tradeoffs to both technical and non-technical audiences, and mentoring engineers to build technical depth and ownership.
- Hands-on with AWS and modern cloud-native tooling (Kubernetes, Docker, Kafka) across microservices and distributed systems delivered to production customers.

SKILLS

Programming Languages: Java, Python, Go, C#, JavaScript

Testing & Quality: JUnit, automated/unit/integration testing, code coverage, code reviews, CI/CD (Jenkins, Bitbucket, GitHub Actions)

Frameworks / Platform: Spring Boot, Django, Flask, Node.js, NestJS, React, .NET Core, Kubernetes, Docker, Kafka, RabbitMQ, ELK Stack, Prometheus, Grafana

Software Development: OOP, RESTful API Design, gRPC, Microservices Architecture, Data Modeling, Design Patterns, ORM, SOA, Strangler Fig Pattern, Event-Driven Architecture

Databases: SQL, NoSQL, PostgreSQL, MySQL, MongoDB, Cassandra, DynamoDB, Redis

Cloud & DevOps: AWS (EC2, EKS, RDS, S3, Lambda), Azure, GCP, CI/CD, IaC, GitOps, Observability, Monitoring & Alerting, Containerization, Orchestration

AI / Data Tooling: Elasticsearch (Apache Lucene), PyTorch, Hugging Face Transformers, Pandas, NumPy, data pipeline design & automation

Collaboration: Jira, Confluence, Git/GitHub/GitLab/Bitbucket, Agile/Scrum

WORK EXPERIENCE

Software Engineer — Ascensus

Mar 2024 – Present

- Architect and build distributed, cloud-native backend services for the Identity & Trust Platform, clarifying ambiguous requirements and partnering with product stakeholders to define the technical roadmap.
- Design the data models and scalable APIs serving as the single system of record for customers, entities, and users across the platform.
- Lead migration of identity and customer-lifecycle logic from legacy systems into resilient, observable, platform-first microservices.
- Build core platform capabilities including a unified entitlements layer and shared compliance-context components.
- Partner with Product, Compliance, and engineering teams to standardize integration patterns and service contracts, communicating design decisions to both technical and non-technical stakeholders.
- Conduct code reviews and mentor junior engineers, driving test coverage and operational best practices across the team.

Tech: Java, Spring Boot, Python, REST, gRPC, AWS, Kafka, Kubernetes, Docker, PostgreSQL, JUnit, Git, Jenkins

Software Engineer Intern, AI/ML — Guidepoint

Oct 2023 – Jan 2024

- Built an AI-based matching system (PyTorch, CUDA) to improve client–advisor pairings, applying Hugging Face Transformers for NLP — improving match relevance 25% and reaching 90% accuracy.
- Built large-scale data preprocessing pipelines (Pandas, NumPy) to prepare high-quality training data for model development, cutting processing time 30%.
- Used Elasticsearch (built on Apache Lucene) for large-scale data indexing and retrieval to support real-time matching.
- Containerized and deployed services with Docker and Kubernetes in a distributed environment, improving scalability 50%.
- Built CI/CD pipelines (Bitbucket) for continuous integration and delivery.

Tech: PyTorch, CUDA, Hugging Face Transformers, Elasticsearch, Pandas, NumPy, Docker, Kubernetes, Python, REST APIs

Software Engineer — Wells Fargo

Apr 2021 – Aug 2022

- Designed and built backend services for financial applications successfully delivered to production, focused on high availability and scalable microservices in a regulated enterprise environment.
- Implemented REST API endpoints with clear, consistent service contracts for internal and external integrations.
- Refactored legacy codebase components into cloud-native architectures, reducing technical debt and improving maintainability.
- Built data-processing and security components meeting compliance and security standards.
- Championed code reviews and automated testing practices, measuring test coverage and improving diagnostics; mentored junior engineers.

Tech: Java, Spring Framework, Python, SQL, Oracle, AWS, Docker, Kubernetes, Jenkins, JUnit

Software Engineer, Production Support — Infosys

Aug 2019 – Mar 2021

- Diagnosed and resolved production incidents and performance bottlenecks for key business applications, minimizing service disruption.
- Partnered with development teams to build permanent fixes for recurring issues and authored troubleshooting/runbook documentation for technical and non-technical audiences.
- Conducted root cause analysis and monitored system performance to proactively prevent future incidents.
- Supported deployment of application patches and new components; managed support tickets against SLAs.

Tech: Java, J2EE, SQL, MySQL, Shell Scripting, Linux, Oracle, Control-M, Nagios

PROJECTS

PipelineOps — Polyglot job-monitoring & alerting platform

- Built a dead-man's-switch monitoring platform for scheduled/batch jobs across a React/TypeScript dashboard, a Django REST Framework API, and a Go/Gin heartbeat-ingestion service, using each language where it fit rather than as a single stack.
- Implemented concurrent heartbeat ingestion in Go with context timeouts, structured logging, and Prometheus metrics; computed job health status server-side with explicit handling for a time-bound "unknown" state.
- Replaced token-based auth stored in browser localStorage with httpOnly session cookies and CSRF protection across the Django backend and React frontend, closing an XSS token-theft vector; added CSRF-enforcement tests verifying writes are rejected without a valid token.
- Established CI (GitHub Actions) running lint and automated tests across all three services on every push, plus Docker Compose for local orchestration and Kubernetes manifests for EKS deployment.

Tech: React, TypeScript, Django, Django REST Framework, Go, Gin, PostgreSQL, Redis, Celery, Docker, Kubernetes, GitHub Actions, AWS

SplitEasy — Expense-splitting app with automated settle-up

- Built a group expense-splitting app (NestJS/Prisma API, React/TypeScript frontend) modeling groups, recurring expenses, invites, and multi-type expense splits (equal, exact, percentage).

- Implemented a minimum-cash-flow settle-up algorithm from scratch using a binary max-heap to greedily match the largest creditor and debtor each round, reducing a group's outstanding balances to a small set of suggested payments.
- Designed a refresh-token auth flow: short-lived JWT access tokens held in memory, backed by a hashed refresh token in an httpOnly cookie, so a database leak or XSS payload can't yield a usable long-lived credential.
- Wrote a full end-to-end test suite (Supertest against a real Postgres instance in CI) covering registration, invites, expense creation, and balance/settle-up computation, alongside unit tests for every backend service and frontend auth/API layer.

Tech: React, TypeScript, NestJS, Prisma, PostgreSQL, Jest, Vitest, React Testing Library, GitHub Actions, Docker

CERTIFICATIONS

- AWS Certified Solutions Architect – Associate
- Certified Kubernetes Application Developer (CKAD)
- Oracle Certified Professional, Java SE Programmer

EDUCATION

Master of Science, Computer Science — University of Massachusetts, Boston
Bachelor of Engineering, Computer Engineering — University of Mumbai

Sep 2022 – May 2024
Aug 2017 – Jun 2021