

Srikanth Guduri

Senior Software Engineer | Java Developer | Full Stack Engineer

913 742 9655 | sguduri997@gmail.com | linkedin.com/in/guds12 | github.com/guds12

PROFESSIONAL SUMMARY

- 5+ years of experience building enterprise backend systems using Java, Spring, and Spring Boot across financial services, retail, and government platforms.
- Skilled in designing and building RESTful APIs and microservices for high traffic domains including licensing, permits, inventory, and product catalog management.
- Experienced in building Apache Spark data processing jobs to handle large scale batch workloads that outgrew traditional SQL based scheduled jobs.
- Proficient with Apache Kafka, SNS, and SQS for event driven architecture, building producer/consumer pipelines that decouple services and support real time data flow.
- Strong SQL and RDBMS background, writing complex queries and stored procedures across Oracle, PostgreSQL, SQL Server, and MySQL, with a focus on schema design over quick fixes.
- Hands on experience with AWS services including S3, EC2, Lambda, ECS, and EKS for both application hosting and data pipeline storage/landing zones, provisioned through Terraform.
- Experience with Databricks for large scale data transformation and consolidation pipelines, building on core Spark skills.
- Built Micro Frontend architectures and custom React component libraries, and works with Node.js and WebSockets for real time, interactive features.
- Applied observability practices using Grafana, Splunk, Dynatrace, Datadog, CloudWatch, and OpenTelemetry, tracing issues across distributed systems rather than guessing from a single log file.
- Comfortable owning backend features end to end, from data modeling through API design and delivery.
- Skilled in containerization and orchestration using Docker and Kubernetes for deploying Spring Boot services and Spark jobs.
- Experienced in writing unit and integration tests using JUnit 5, Mockito, and REST Assured, with SonarQube driven code quality reviews in CI/CD pipelines.
- Practiced in Agile/Scrum methodologies, including sprint ceremonies, production support rotations, and methodical log based debugging.
- Mentored engineers and collaborated cross functionally with business and technical stakeholders on secure, scalable system design.
- Experienced in using enterprise-approved AI-assisted engineering tools to support code development, refactoring, code reviews, test creation, troubleshooting, and root-cause analysis.
- Applied responsible AI practices by validating AI-generated outputs for correctness, performance, security, and data-sensitivity concerns before incorporating them into production solutions.
- Mentored engineers on secure coding, peer reviews, automated testing, AI-output validation, and the reuse of effective engineering patterns across teams.
- Experienced in evaluating application architectures and emerging technologies for security, scalability, integration feasibility, and alignment with existing enterprise systems.
- Completed the Microsoft **AI-103** certification, demonstrating knowledge of building, managing, and deploying AI applications and agents on Azure.

TECHNICAL SKILLS

Languages	Java 8/11/17, SQL, Python
Backend	Spring Boot, Spring MVC, Spring Security, Spring Data JPA, REST APIs, Microservices, Node.js, Redis, Akamai
Frontend	React.js, Angular, TypeScript, JavaScript, Micro Frontends, WebSockets
Data Processing	Apache Spark, Databricks
Messaging	Apache Kafka, AWS SNS, AWS SQS, Event Driven Architecture
Databases	Oracle, PostgreSQL, SQL Server, MySQL, RDBMS Design
Cloud and IaC	AWS (S3, EC2, Lambda, ECS, EKS, CloudWatch), Microsoft Azure, Terraform
Observability	Grafana, Splunk, Dynatrace, Datadog, OpenTelemetry
CI/CD and DevOps	Git, Jenkins, GitHub Actions, Maven, Docker
AI	AI-Assisted Coding, Code Review, Refactoring, Test Acceleration, Root-Cause Analysis, AI-Output Validation, Responsible AI, Secure AI Usage
Domain	Commercial Banking, Consumer Banking, Regulatory Compliance
Testing	JUnit 5, Mockito, Integration Testing
Process	Agile and Scrum, JIRA, Code Reviews

PROFESSIONAL EXPERIENCE

Senior Software Engineer

Aug 2023 to Present

NYC Office of Technology and Innovation, New York, NY

- Developed Java Spring Boot backend services for the NYC Business Portal, building REST APIs for the licensing, permits, and inspections modules.
- Built Kafka based event pipelines between portal microservices, moving citizen workflow events like application submitted and status changed without tightly coupling services.
- Wrote complex SQL queries and stored procedures against PostgreSQL for reporting features, restructuring a couple that had been scanning far more data than needed.
- Used AWS S3 for document storage and as a landing zone for Spark job outputs before they were loaded into downstream reporting tables.
- Used Databricks for a data transformation pipeline that consolidated records from multiple agency source systems into a single cleaned dataset.
- Applied RDBMS design principles when modelling new citizen data domains, normalizing where it made sense and denormalizing deliberately where read performance mattered more.
- Deployed Spring Boot services to AWS EC2 and Lambda depending on whether a service needed to run continuously or could be triggered on demand, provisioning the underlying infrastructure with Terraform.
- Built React and TypeScript components for the citizen facing sections of the portal, working closely with the backend REST APIs to deliver a complete end to end feature.
- Split part of the portal frontend into independently deployable Micro Frontends, letting different teams release their sections without coordinating a single monolithic build.
- Used enterprise-approved AI-assisted development tools to accelerate Java code refactoring, unit-test creation, technical documentation, and troubleshooting while independently validating generated outputs.
- Established consistent validation practices for AI-assisted development, including peer reviews, automated testing, SonarQube analysis, security checks, and performance verification.

- Applied responsible AI practices by avoiding sensitive data in prompts, reviewing generated code for security risks, and ensuring AI-assisted solutions followed organizational security and resiliency standards.
- Used AI-assisted log and code analysis with Splunk and CloudWatch to support root-cause investigations and reduce the time required to diagnose application failures.
- Identified recurring service and data-processing failures and improved operational stability through automated monitoring, alerts, validation checks, and reusable troubleshooting procedures.
- Led technical discussions with internal engineering teams to evaluate architectural designs, technology compatibility, security requirements, and integration with existing systems.
- Shared reusable AI-assisted development, testing, and troubleshooting patterns with engineers and supported team adoption through technical guidance and knowledge-sharing sessions.
- Monitored portal services using AWS CloudWatch, setting up alarms for Lambda error rates and EC2 resource utilization.
- Wrote unit tests with JUnit 5 and Mockito for new service logic, and reviewed SonarQube findings before merging into the shared branch.
- Documented REST endpoints with Swagger so frontend and integration teams could work from accurate contracts instead of checking in with the backend team constantly.
- Troubleshooted a Spark job that was failing intermittently under larger data volumes, tracing it to a memory configuration that had not been adjusted since the job was first written.
- Participated in Agile sprint ceremonies including planning, standups, and retrospectives, and kept stakeholders informed of progress and risk.

Tech Stack: Java 17, Spring Boot, React, TypeScript, Micro Frontends, Apache Spark, Databricks, Kafka, PostgreSQL, AWS S3, AWS EC2, AWS Lambda, AWS CloudWatch, Terraform, JUnit, Mockito, JIRA

Software Development Engineer

Feb 2021 to July 2022

Michelin, Hyderabad, India

- Developed Java Spring Boot microservices for Michelin's retail platform services, covering product catalog, inventory, and customer personalization domains.
- Wrote Spark jobs for retail analytics workloads, processing product and pricing datasets that had grown too large for the team's existing relational reporting setup.
- Implemented Kafka producer and consumer pipelines to process real time inventory updates from warehouse systems, decoupling producers from slower downstream consumers.
- Used AWS SNS and SQS for fan out notification patterns and decoupled task queues where a full Kafka topic would have been more infrastructure than the use case needed.
- Wrote and optimized SQL queries for reporting and analytics features against Oracle, restructuring a slow product search query that was scanning far more of the catalog than necessary.
- Used AWS S3 for retail platform document storage and as an intermediate landing zone for Spark job inputs and outputs.
- Automated recurring monitoring and quality checks using Jenkins, SonarQube, Splunk, Datadog, and Dynatrace, improving early detection of application and deployment issues.
- Coached engineers on service architecture, secure development, code-review standards, automated testing, observability, and production troubleshooting.
- Evaluated technical approaches for Kafka, SNS, SQS, ECS, and EKS based on scalability, resiliency, operational complexity, and integration requirements
- Used Databricks for a batch pipeline that scored product data against a set of quality checks before it went live on the storefront.

- Applied RDBMS design practices when modelling new inventory and pricing schemas, balancing normalization against the query patterns the retail platform actually needed.
- Built Angular components for internal merchandising tools, consuming the same Spring Boot REST APIs used by the broader retail platform.
- Built a Node.js service layer with WebSocket support for a real time inventory status feature, pushing live stock updates to the storefront instead of relying on polling.
- Contributed to a Micro Frontend setup for parts of the customer facing storefront, allowing separate teams to own and deploy their sections independently.
- Deployed containerized services using Docker and Kubernetes on AWS ECS and EKS, choosing between them depending on how much orchestration control a given service actually needed.
- Set up observability using Splunk, Dynatrace, Grafana, and Datadog dashboards, and instrumented services with OpenTelemetry for distributed tracing across the microservices handling a single checkout request.
- Wrote tests across the stack, including JUnit and Mockito for unit coverage and REST Assured for API level checks, with SonarQube tracking coverage in the Jenkins pipeline.
- Participated in code reviews and Agile ceremonies, and helped onboard a couple of new team members to the service architecture.

Tech Stack: Java, Spring Boot, Angular, Node.js, WebSockets, Micro Frontends, Apache Spark, Databricks, Kafka, AWS SNS, AWS SQS, Oracle, AWS S3, AWS ECS, AWS EKS, Docker, Kubernetes, Splunk, Dynatrace, Grafana, Datadog, OpenTelemetry, Jenkins, JUnit, Mockito

Software Engineer

Dec 2019 to Jan 2021

Redvue Technologies Pvt Ltd, Chennai, India

- Developed Java Spring Boot microservices for Michelin's retail platform services, covering product catalog, inventory, and customer personalization domains.
- Developed and maintained Java-based backend applications using Java, J2EE, Spring Boot, and Spring MVC.
- Built and integrated REST APIs to support application features such as account management, data processing, and transaction-related workflows.
- Wrote SQL queries and stored procedures using Oracle and SQL Server to retrieve, update, and manage application data.
- Implemented API security using Spring Security, including authentication and role-based authorization.
- Worked with AngularJS and jQuery to develop and maintain web application screens and integrate frontend components with backend REST APIs.
- Used Maven for project builds and Git for source code management and version control.
- Assisted with Jenkins CI/CD pipelines for automated application builds, testing, and deployments.
- Created JUnit unit tests to validate application functionality and identify defects.
- Participated in Agile development practices, including sprint planning, daily stand-ups, code reviews, and defect resolution.
- Collaborated with developers, QA engineers, and business teams to understand requirements and deliver application enhancements.

EDUCATION

M.S. Computer Science, University of Central Missouri, Warrensburg, MO, Dec 2023

B.Tech. Computer Science, KL Deemed University, Vaddeswaram, AP, India, May 2020

CERTIFICATIONS

- AWS Certified Developer – Associate (**DVA-C02**)
- Microsoft Certified: Azure AI Apps and Agents Developer Associate (**AI-103**)