

MANOJ SANAGAPALLI

SOFTWARE ENGINEER 2 • BACKEND & DISTRIBUTED SYSTEMS • AI AGENTS

Bengaluru, India | +91 9491130256 | sanagapallimanoj@gmail.com | linkedin.com/in/manojsmkg/

01 PROFILE

Software Engineer with 6+ years of experience building scalable backend systems using Java, Spring Boot, Kafka, Spark, Numaflow, and microservices. Experienced in distributed systems, production reliability, workflow automation, data-intensive platforms, and AI-agent integrations, with end-to-end ownership from system design through production delivery.

02 CORE SKILLS

- Languages: Java, Scala, Python, C/C++
- Backend & Architecture: Spring Boot, JPA, REST APIs, Microservices, Hexagonal Architecture, Event-Driven Architecture
- Distributed & Data: Apache Kafka, Numaflow, Apache Spark, Spark Streaming, Hadoop, Hive
- Databases: Teradata, MySQL, DB2, Sybase
- AI & Integration: MCP, AI Agents, OAuth, REST integrations
- Cloud & DevOps: Docker, Git, Jenkins, CI/CD, OpenTelemetry, HPA
- Engineering: System Design, HLD, LLD, Database Modelling, Performance Optimization, Reliability

03 PROFESSIONAL EXPERIENCE

INTUIT | Software Engineer 2

Apr 2025 – Present • Bengaluru, India

- AI-Agent Bulk Import: Designed a reusable ingest → preview → submit → status architecture exposed as MCP tools, then extended the platform across multiple QuickBooks Online entity types; owned the path from backend data pipeline to callable Omni-agent skill.
- Customer Bulk Update: Built a full-stack customer bulk-update wizard with match/diff/create upsert semantics against the canonical QBO schema, including ambiguous-match handling and user-review staging; validated the ~6,200-line implementation with live performance testing.
- Cost Migration & Data Platform: Shipped a unified Convergence migration engine and a new lifecycle state without forking the existing state machine; combined report fetching to halve downstream QBO API calls and redesigned large-company budget purge into bounded chunked deletion tasks.
- Event-Driven Migration: Migrated billing-event consumption from Kafka to a real-time Numaflow pipeline using a hexagonal dual-mode architecture for zero-risk cutover; production processing includes 1,296+ billing events and 379+ customer migration jobs, with 4,125+ migration jobs triggered over a recent 3-month window.
- Production Reliability: Resolved connection-pool exhaustion, startup binding failures across all pods in 3 availability zones, hot-table query-plan issues, and numeric-precision defects; improved observability and drove RCA/preventive actions including compatibility gating and health-check isolation.
- Workflow Automation & Scale: Delivered Bill Payment Approval Workflow end-to-end across Approval Service, Workflow Automation Service, and AppConnect; 5,744+ approval actions across 605+ workflow definitions with zero attributed production incidents. Also led tax-season resiliency work across HPA, HikariCP tuning, async processing, and OpenTelemetry.

MORGAN STANLEY | Software Engineer III

Jan 2023 – Apr 2025 • Bengaluru, India

- Real-Time Data Processing: Architected Spark Streaming + Kafka processing for 30,000 financial adjustments, reducing processing time from 30 minutes to 2 minutes and improving data consistency and operational efficiency by 93%.
- High-Scale ETL: Designed Apache Spark ETL workflows supporting allocation of \$3.5B in monthly P&L, improving data transformation efficiency for finance operations.
- Tax Automation: Engineered a full-stack Angular + Java Spring Boot solution that eliminated \$2B in high-risk manual entries annually, saving 16 hours/month and reducing data-entry errors by 99.9%.
- Microservices Architecture: Led migration toward modular microservices, enabling independent scaling, reducing deployment times by 70%, and improving overall system reliability.
- Delivery & Analytics: Led requirements-to-production delivery for critical financial software and created Tableau analytics/reporting used by finance stakeholders for data-driven decisions.

MORGAN STANLEY | Software Engineer II

Jan 2022 – Jan 2023 • Mumbai, India

- Reference Data Platform: Built a full-stack static reference-data management tool using Angular and Spring Boot REST APIs, decoupling dynamic reference-data changes from application code.
- Config-Driven APIs: Developed Spring Boot + Spring JPA APIs for CRUD operations on Teradata, simplifying onboarding of new reference-data tables and improving development efficiency.

MORGAN STANLEY | Software Engineer I

Aug 2020 – Jan 2022 • Mumbai, India

- **P&L; Calculator:** Developed a high-performance Apache Spark calculator processing terabytes of financial data daily and improved unclaimed-balance computation accuracy by 99%.
- **Developer Productivity:** Contributed to DevOps Hackdays addressing Jenkins, build, version-control, and release-management pain points to streamline engineering workflows.

04 ACHIEVEMENTS

- Recognized for exceptional contributions and awarded for exceptional delivery within the India Technology Division at Morgan Stanley.
- Received back-to-back promotions over two years and converted from intern to full-time employment, reflecting increasing scope and responsibility.

05 EDUCATION

SASTRA Deemed University — BTech, Computer Science

Thanjavur, India • GPA: 8.5 / 10.0 • May 2020

Coursework: Data Structures, Algorithms, Operating Systems, Database Management, Artificial Intelligence, Machine Learning, Data Mining & Data Warehousing

06 RESEARCH & PROJECTS

- **Speech-Assisted Systems:** Co-authored the research paper “Multi-view Representation Based Speech Assisted System for People with Neurological Disorders” and developed a Python-based machine-learning model supporting the study.
- **Mess Menu Prediction:** Built a menu prediction tool using student preference data collected through Google Forms; applied the Apriori algorithm to identify frequent item sets and popular menu choices.