

# SIDDHARTH PANDA

## Senior Software Engineer | Distributed Systems & AI Infrastructure

sidkiit12@gmail.com • +91 7749821508 • Bangalore, India • [LinkedIn](#) • [GitHub](#)

### PROFILE

Senior Software Engineer building distributed systems, AI infrastructure, and developer-facing platforms. Experience scaling to 1T+ daily events at LinkedIn across Kafka, Flink, Go, Kubernetes, and MCP. Interested in high-ownership startup roles building foundational products end to end.

### EXPERIENCE

#### LinkedIn

##### Senior Software Engineer

04/2026 - Present

- Created a Managed Event Enrichment Platform using configuration-driven Flink processes for Kafka events, reducing onboarding time from ~2 days to ~1 hour.
- Built a Tracking MCP server to automate event-schema and pipeline onboarding in agentic workflows.
- Designed Cost-to-Track, attributing compute, network, and storage costs across Kafka/Flink service tiers; modeled high-cardinality metrics in Apache Pinot with Star-Tree indexes for sub-100ms P95 queries over billions of records.
- Integrated high-throughput streaming pipelines with the Realtime Feed Recommendation Engine to provide real-time contextual signals.

##### Software Engineer

05/2024 - 03/2026

- Developed and maintained telemetry infrastructure supporting 1T+ daily events, improving reliability and reducing telemetry regression detection from days to minutes.
- Led migration of stream-processing systems from Apache Samza to Flink, achieving 97% performance parity across 12 parity jobs.
- Collaborated with HDFS, Kafka, and Schema Registry teams to deliver critical infrastructure solutions and improve system robustness.

#### Grab

##### Senior Software Engineer

03/2023 - 04/2024

- Increased lending repayment rate by 1.2 percentage points by building a Kafka-based notification system and partnering with PMs on product pain points.
- Introduced a line-of-credit mechanism for DAX partners by restructuring existing loans, increasing disbursement by ~\$2M.

##### Software Engineer

08/2021 - 03/2023

- Re-architected the financial auto-deduction flow, reducing batch execution time from 1.5 hours to 20 minutes.
- Set up CI/CD pipelines for post-borrowing microservices, cutting QA regression cycle time from 3 days to 1.5 days.

#### HighRadius

##### Software Engineer Intern

07/2020 - 04/2021

- Automated sensitive PII encryption, reducing environment cloning time from 2 hours to 5 minutes.

#### Khosla Labs

##### Software Engineer & ML Intern

05/2019 - 07/2019

- Built document-capture models, OCR pipelines, and computer-vision image-processing workflows using OpenCV.

### PROJECTS

#### DumbTF - Deep Learning Framework

- Built an autograd and computational-graph deep learning framework from scratch in Python using only NumPy; benchmarked convergence on MNIST, Boston Housing, and synthetic regression datasets.

#### Document Detection & Scanner - OpenCV

- Built a document-detection pipeline using grayscale conversion, thresholding, Canny edge detection, contour analysis, and geometric filtering; integrated the approach into an Android scanning application.

### TECHNICAL SKILLS

**Languages:** Go, Java, Python, C++, Scala, Rust, Bash, SQL

**Streaming & Data:** Apache Kafka, Apache Flink, Apache Spark, Apache Samza, Airflow, HDFS, Elasticsearch, MySQL, Apache Pinot

**Infrastructure:** Distributed Systems, Event-Driven Architecture, Kubernetes, Docker, CI/CD, REST APIs

**AI & Developer Tools:** MCP, AI-assisted Development, GitHub Copilot, Cursor, Claude Code

### EDUCATION

Kalinga Institute of Industrial Technology - B.Tech, Computer Science & Engineering | 2017 - 2021

GPA: 9.17/10