

Ravi Prakash

Professional with 7+ years of full time experience in Backend Software Development (Golang, Python, Java), MLOps, IOT across several product-based startups. Ability to hack around multiple technologies and build highly scalable, low latency, concurrent, distributed systems, RESTful services

+91-9868576600

prakashravip1@gmail.com

[LinkedIn](#) | [Portfolio](#) | [Github](#) | [Leetcode](#) | [Medium](#) | [StackOverflow](#) | [Designs](#)

TECHNICAL SKILLS

Programming Languages: Golang, Python, Java

Databases: Cassandra, GCP Bigtable, Neo4J, DynamoDB, Apache Pinot, MongoDB, Elasticsearch, Kafka, AWS SQS, SNS, RabbitMQ, MySQL, PostgreSQL, GCP BigQuery, ClickHouse, Redis, Aerospike, Solr, Delta lake

Frameworks and Libraries: Spring-Boot, Go-gin, go-mux, gorm, Django, Flask, FastAPI, Sanic, Pandas, Keras, Computer Vision, Asyncio, Terraform, Apache Spark

Monitoring and Observability: Prometheus, Datadog, Grafana, Kibana, ELK, Google Cloud Trace, Sumologic

EDUCATION

[**NSUT \(formerly NSIT, University of Delhi\)**](#), New Delhi — *Bachelor in Engineering*

2014 - 2018

Branch - Computer Engineering, CGPA 8

EXPERIENCE

AiStrike, Remote — Senior Engineer

Dec 2024 - Present

- **Architected** and implemented enterprise-scale microservices using Spring Boot and OpenJDK 17, integrating **PostgreSQL, Neo4j**, and **MongoDB** in a **Kubernetes** environment, ensuring high availability and scalability.
- Led development of mission-critical security systems including a Risk Booster platform with advanced **Neo4j** query capabilities and an AI-powered Compliance Questionnaire system leveraging **AWS S3**
- Engineered automated **AWS infrastructure** solutions including multi-account onboarding and cross-regional monitoring systems, improving operational efficiency and resource management using **Python(FastAPI)**
- Implemented comprehensive security enhancements across **Docker** containers and developed a centralized vulnerability management platform integrating multiple scanning sources
- Optimized system performance through **Redis** caching implementation and **Kafka**-based real-time data synchronization, significantly improving response times and indexing client data in **Apache Solr** cluster.
- **Skills: Python, Golang, Java, Neo4J, PostgreSQL, MongoDB, Solr, Kubernetes, Redis, Kafka, AWS**

Brightedge, Remote — Software Engineer 3

Aug 2023 - Dec 2024

- **Led the initiative** in page based recommendation team by maintaining and enhancing backend systems thereby improving recommendation accuracy from 92% to 98% for global and corporate accounts.
- Developed and optimized internal admin tools and backend systems, including a centralized data source for local businesses and a public API for accessing local business data.
- **Architected and led a team of 4 Junior Dev** implementing end to end features for monitoring and syndicating business listings across Google, Apple, and Bing Maps, including holiday hours, service area management, centralised display of business insights and review management system for new businesses on the Brightedge Local business management platform. Incorporated trial features, businesses onboarding flow for seamless integration.
- **Skills: Python, Golang, PostgreSQL, Bigquery, Apache Solr, Docker, Kubernetes, Redis, GCP, Kafka, AWS**

BatterySmart, Gurugram — SDE-3

Jan 2023 - June 2023

- Develop EV battery's IOT infrastructure by collecting IOT events from multiple battery vendors and storing them in a **Apache Pinot** realtime table. It is able to handle scale of over 70 qps and manage TBs of data storage and power real time analytics. **Tech used - Golang, Kafka, Avro, Apache Pinot, Prometheus, Grafana, Redis, Trino, Redash, Kubernetes**
- Written a kafka producer to send **Avro** serialized electric battery IOT event data for a vendor. Tech used - **Python, Avro, Kafka, Kubernetes**
- Developed an Insights service which captures IOT event for every EV Battery swap transaction and stores it in a **Apache Pinot** realtime table. **Tech used - Golang, Kafka, Pinot, Kubernetes**
- Developed showing live EV battery location based on IOT data. **Tech used: NodeJS, Redis, Kubernetes, Mysql, Redis Geolocation**

Tekion Corp, Bengaluru — ML Engineer

July 2021 - Aug 2022

- Driving design and development of **distributed batch** and **streaming ETL** processing pipeline
 - Responsible for setting up a scalable, cheap, ACID compliant **delta lake** storage over **AWS S3** for storing processed raw facts known as features and batch prediction results of ML models, replacing **PostgreSQL**. [Blog](#)
 - Responsible for development and maintenance of internal python clients and tools packages using **poetry, pip**.
 - Optimized top k recommendation API latency by doing batch **cosine similarity** calculation and **serializing high dimension feature vectors** in **Apache Parquet** and storing it in an **FAISS embedding store**.
 - Created a bare SQL parser using **pyarrow** for ML data pipeline to facilitate direct read of raw data from delta lake warehouse. It supports query filtering over **partitioned and non-partitioned delta lake columns**, thereby reducing the need and cost of querying a managed query engine like **AWS Athena**
 - Architected a system for merging duplicate sales leads using **Dedupe, Spark, Delta lake, Kafka, AWS Athena**.
 - Splitting **sales chat app** built using **RASA** framework into containerized (core and ACTION) microservice to facilitate communication among them and integrated our observability tools and dynamic model loading.
 - Improved ML serving APIs performance by converting all I/O bound tasks into asynchronous calls. Migrated from **flask** to **FastAPI** and added parallelism through **uvicorn** workers and added API observability.
- Skills - **opentelemetry, elastic-apm, Elasticsearch, Kibana**

magicpin, Gurugram

Tech Lead

Aug 2022 - Jan 2023

- **Led organisation wide Observability Initiatives**
- Deployed **Prometheus** in GKE for monitoring and integrated clients in **Go, Java, and Python**.
- Optimized metric ingestion with relabeling filters.
- Designed serving architecture for **Google Vertex AI** model using **FastAPI** and **Bigtable**.
- Built pipelines for feature calculation, real-time predictions, and data sync between **BigQuery** and **Bigtable**.
- Implemented **Kubernetes** jobs for model scoring and established **CI/CD** workflows for deployment.

Software Engineer

June 2018 - Dec 2020

- Responsible for continuously evaluating recent technologies and driving architecture and design decisions on various critical backend systems (**Social Media Feed, Order Management System, Search, Transaction**) and ownership to run and maintain **GCP** cloud infrastructure. **Skills: Apache Solr, Cassandra, Neo4J, Elasticsearch**
- Performed **dataset collection, feature engineering**, training using transfer learning on **InceptionResnetV2** and developed model inference using **RESTful flask service** with **Nginx** Load Balancer for [selfie classification](#)
- Migrated 21 TB data from **21 node Cassandra cluster** to **Google Cloud Bigtable** cluster with redesigned schema ensuring **row key entropy**.
- Migrated a tightly coupled **Java-MySQL monolith** backend stack to decoupled, low latency, **Golang microservice** with **denormalized MySQL schema** thereby improving availability and reducing infra costs.
- Designed and developed a new low latency, highly available **reverse-proxy layer** using **Golang (gin)** between **backend microservices** and App/Website with proper **Authentication** and **Authorisation** mechanism.
- Designed and developed **concurrent 5 layer trending pipeline (Kafka, Golang, Elasticsearch, Cassandra)** which generates scores accumulated over days with decaying functions for various social entity in our app like posts, hashtags, merchants, users, friends and using **worker-pool, pipeline, barrier concurrent** design patterns
- Designed and implement **Kafka** consumer and **Golang (mux)** API Flow for projects like Merchant Rating calculation based on Bayes estimation, merchant complaints and review flow and Hashtag Feed and Campaign Feed and selfie and hashtag approval flows, groupbuys, merchant and user feed.