

Ravi Prakash

Professional with 6.5 years of full time experience in Backend Software Development (Golang, Python), MLOps, IOT across several product-based startups. Ability to hack around multiple technologies and build highly scalable, low latency, concurrent, distributed systems, RESTful services

9868576600

prakashravip1@gmail.com

[Github](#)

[LinkedIn](#)

[Blog](#)

EXPERIENCE

Brightedge — *Software Engineer 3*

August 2023 - Present

- Worked as part of 3 member team maintaining, developing backend and hadoop clusters of webpage SEO recommendations.
- Performed deep dive into accounts page based recommendations and improved accuracy from 92% to 98% for global and corporate accounts.
- Optimised and reduced processing time of scripts from over 5 hours to under 45 minutes by doing DB optimisation, implementing code flow efficiency and introducing helpful thresholds
- Optimised Google, Bing and Apple location flows like service areas, service items, reviews, location and campaign posts, operating hours
- Managed the successful integration of location-based services for more than 500 businesses into Apple Maps while ensuring compliance with all protocols
- Developed centralized data source for local businesses in maps as entity which then syndicated to Google, Apple and Bing Maps.
- Initiated and developed the flow of fetching and storing google and bing insights like (clicks, views, impressions from desktop, mobile sources) from business APIs and created APIs on top of it for UI Dashboard
- Exposed a public API for businesses to read their local business data present with Brightedge, so that they can display it on their websites.
- Designed and developed flow to store Google Maps reviews and calculate average ratings and showing only new reviews to the users.
- Designed several processes to introduce special holiday hours with labels for businesses and then syndicate it to bing, apple and google maps.

BatterySmart, Gurugram — *Software Engineer 3*

January 2023 - June 2023

- Developed an EV battery's IOT infrastructure by collecting IOT events from multiple battery vendors and storing them in a realtime table. It is able to handle scale of over 70 qps and manage TBs of data storage and power real time analytics.
- 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.

magicpin, Gurugram — *Tech Lead*

August 2022 - Jan 2023

- Took lead in setting up organisation level observability using Prometheus and Grafana.
 - Did Prometheus client integration into Go, Java, Python microservices.
 - Reduced organisation wise metric samples ingestion by applying metric relabeling filters.
- Designed and productionised a real time ML prediction and serving architecture for Google Vertex AI trained statistical ML model.

Tekion Corp, Bengaluru — *ML Engineer*

July 2021 - Aug 2022

- Driving designing and development of distributed batch and streaming ETL processing pipeline, offline and online ML model serving architecture.
- Responsible for setting up a scalable, cheap, ACID compliant delta lake storage

TECHNICAL SKILLS

LANGUAGES

Golang, Python, C++, Java

DATABASES

Delta Lake, Data Lake

Cassandra, Bigtable, DynamoDB, Apache Pinot

MongoDB, ElasticSearch,

Kafka, SQS, RabbitMQ

MySQL, Postgres, Bigquery

Redis, Aerospike

Neo4J, Treno, Redash

FRAMEWORKS

Go-gin, go-mux, gorm, Pandas, Django, Flask, FastAPI, Avro, Keras, Tensorflow

AWS, Azure, GCP

Docker, Kubernetes

Prometheus, Datadog, Grafana, Kibana

REST, gRPC, JSON, GraphQL

over S3 for storing processed raw facts known as features and batch prediction results of ML models, replacing postgres.

- Responsible for development and maintenance of several internal python clients and packages used by ML Operations and data science team.
- Optimized top k recommendation API latency by doing batch cosine similarity calculation and serializing high dimension feature vectors in parquet and storing it in an in-memory based fast storage.
- Created a bare SQL parser over 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 columns
- Developed a scalable architecture for sharing duplicate data merge pipeline results using spark, delta lake, Kafka, and AWS Athena.
- Splitting sales chat app built using RASA framework into containerized (core and ACTION) microservice to facilitate communication among them..
- Added concurrency to ML serving APIs by converting all I/O bound tasks into asynchronous calls and added parallelism using multiple uvicorn workers.

daloopa, Noida — *Software Engineer III*

December 2020 - June 2021

- Researched and implemented foreign language OCR detection, recognition and translation of financial documents using open source technologies thereby saving costs.
- Designed and implemented a data pipeline to support pattern-based search on financial documents using machine learning, OCR, NLP and heuristics-based search algorithms.
- Improving latency of critical flows through caching, concurrency, multiprocessing and pre-processing.

magicpin, Gurugram — *Software Engineer*

June 2018 - Dec 2020

- Responsible for continuously evaluating recent technologies and driving architecture and design decisions on various critical backend systems and ownership to run and maintain cloud infrastructure.
- Took initiative for dataset collection, feature engineering, training using transfer learning, model inference using RESTful flask service on CNN based deep learning classification model for selfie classification.
- Migrated 21 TB data from 21 node Cassandra cluster to Cloud Bigtable cluster with redesigned schema ensuring row key entropy
- Migrated a tightly coupled java-MySQL monolith backend stack to decoupled, low latency, go service with denormalized MySQL schema thereby improving availability and reducing infra costs.
- Designed and developed a new low latency, highly available reverse-proxy layer between API layer and front-end for several backend systems which resulted in secure and robust backend systems.
- Designed and developed concurrent 5 layer trending pipeline 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 design patterns
- Designed and implemented Consumer and 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
- Designed and developed notification,sms,whatsapp,IVR trigger logic along with an order cleanup service for Order Management Team

EDUCATION

NSIT, New Delhi — *Bachelor in Engineering*

August 2014 - May 2018

Branch - Computer Engineering - CGPA 8

Certificates

[Introduction to Machine Learning in Production](#)

[Go \(Intermediate\)](#)

[Confluent Fundamentals Accreditation](#)

[Programming with Google Go by University of California, Irvine](#)

[Google Cloud Platform Fundamentals: Core Infrastructure](#)

[DeepLearning.AI TensorFlow Developer](#)

AWARDS

Gold Medal - Secured 1st State Rank and 20th International Rank in International Informatics Olympiad (IIO) in 2010.

CBSE Certificate of Merit: Secured grade A1 in all 5 subjects in Xth Board Exam.

Central Sector Scheme for College and University Student

Received scholarship for securing high percentile in JEE Main