

Harsh Verma

harshverma3305@gmail.com | LinkedIn | GitHub | +919654499559

Software engineer focused on data platforms and distributed systems. Built high-throughput backend and data infrastructure in Java, Kotlin, and Python across Apache Flink, Kafka, Spark, Airflow, AWS, Kubernetes, SQL, and Spring Boot.

Experience

Software Development Engineer 2, Expedia Group – Gurgaon, India Jan 2025 – Current

- Architected and deployed an Apache Flink-based streaming pipeline on AWS, scaled throughput from **123 TPS to 75K TPS** by offloading aggregations from MongoDB to Flink's native stateful processing using **RocksDB**, reducing aggregation latency from **300ms to 0.02ms**, and cutting annual cloud costs by **\$48K**.
- Collaboratively developed a personalized review-form platform using traveler and property data, increasing review collection rates by **12.8%** across three major brands.

Software Development Engineer 1, Expedia Group – Gurgaon, India Aug 2023 – Jan 2025

- Built a review sentiment validation platform for **500+ concepts** across **1.5M** properties, processing **12M+** reviews and contributing to an estimated **\$1M+ annual revenue** impact.
- Developed an MVP Content Accuracy Score platform for **100K** properties by analyzing content quality, amenities, and negative review sentiment; improved data accuracy for **46K properties**, impacting **32%** (**\$14.7B**) of GBV and **502M page views** over six months.

Software Development Engineer Intern, Expedia Group – Gurgaon, India May 2022 – Aug 2022

- Migrated **4 Spring Boot** applications from a legacy brand-specific cloud platform to a brand-agnostic cloud platform, reducing costs by **73%**.
- Developed personalized review sorting logic using traveler locale, location, brand preference, and review relevancy to surface more relevant reviews.

Education

IIIT-Delhi, B.Tech in Electronics and Communication Engineering (ECE) Aug 2019 – Aug 2023

- **Coursework:** Data Structures and Algorithms, DBMS, Distributed Systems, Machine Learning, OOP in Java, Computer Vision, Computer Architecture

Projects

Open Source Contribution: Micrometer

- Contributed to open-source Micrometer core ecosystem, developing a feature to enable dynamic tags in Log4j2 counter metrics to enhance real-time application observability

HRV-Mart E-commerce Backend [Documentation Link](#) Jan 2023 – Aug 2023

- Developed a scalable e-commerce backend using microservices, Spring Boot, reactive programming, MongoDB, Kafka, Docker, Kubernetes, and GitHub Actions.
- Built CI/CD deployment on Azure and established testing workflows for maintainable production releases.

Skills and Technologies

Languages: Java, Kotlin, Python

Data & Backend: Apache Flink, Apache Spark, Apache Kafka, Apache Airflow, SQL, Hive, gRPC, REST, Spring Boot

Infrastructure: Claude Code, GitHub Copilot, Gen AI, AWS, Kubernetes, Docker, Azure, Redis, MongoDB, Elasticsearch, Jenkins, GitHub Actions, Maven, Gradle, Git