

Resume

Staff Software Engineer

San Jose, California | ishanshah301190@gmail.com | ishandshah.github.io/portfolio

Summary

Staff Software Engineer with 10+ years of experience designing and delivering distributed systems, cloud-native services, event-driven platforms, real-time data pipelines, billing systems, supply chain platforms, and security analytics infrastructure. Led and contributed to high-throughput billing, real-time inventory, CDC, optimization, and search platforms across PayPal, Nordstrom, and Securonix, with recurring focus on Java, Kafka, AWS, observability, CI/CD, and resilient multi-tenant architecture.

Experience

Dec 2024 to Present

PayPal — Staff Software Engineer

- Early engineer on a cloud-native usage-based billing platform delivered from concept to production in six months.
- Built hybrid billing architecture with REST APIs and Kafka pipelines for 100K+ TPS ingestion and aggregation workloads.
- Implemented multi-tenant, cell-based architecture patterns for merchant isolation, resiliency, and independent scale.
- Developed Java and Spring Boot services across subscription billing, usage aggregation, invoicing, refunds, credit notes, discounts, tax handling, retries, and cost calculation.
- Provisioned AWS and platform infrastructure including Lambda, DynamoDB, S3, RDS, API Gateway, Redis caching, Kubernetes, Docker, and Terraform.
- Integrated OpenTelemetry, Splunk, Datadog, and CloudWatch to improve operational visibility and reduce incident response time.
- Authored merchant-facing API documentation and used AI-assisted engineering tools for code review, design validation, and documentation acceleration.

Mar 2020 to Dec 2024

Nordstrom — Senior Software Engineer

- Enhanced inventory placement systems by factoring speed, cost, and customer proximity using Java, Spring, Kafka, AWS, and Gurobi.
- Collaborated with data science partners on a speed-cost optimizer that replaced distance-only fulfillment logic.
- Built an Optimizely-based experimentation platform to route requests across models and validate strategies through historical testing.
- Implemented marketplace reverse logistics with Java, Spring Boot, AWS, REST APIs, Kafka, and Postgres RDS, supporting revenue recovery opportunities.
- Engineered real-time CDC pipelines with Debezium, Kafka Streams, Kafka state stores, and Avro events to improve inventory availability accuracy.
- Developed unified outbound inventory consumers using Spring Boot, Java, Kafka, Redis, DynamoDB, and Postgres for real-time synchronization.
- Enabled invoice-level inventory analysis and near real-time stock ledger transactions for financial auditing and cost tracking.
- Implemented Kafka producer partitioning strategies to support consumer-level customization and reduce processing lag.

Dec 2016 to Mar 2020

Securonix — Software Engineer

- Worked on a security intelligence platform for real-time intrusion analysis, anomaly detection, and activity trend monitoring from live log streams.
- Developed Kafka and Redis-backed data pipelines for normalization, enrichment, Solr indexing, and full-text analyst search.
- Implemented commands and operators for Spotter DSL, a natural-language search layer based on Solr and HDFS.
- Improved Solr performance by 2x through tuning, document routing, and implicit sharding.
- Implemented recovery and high-availability patterns supporting 95% uptime targets and indexing speeds up to 500K events per second.

Extra Initiatives

- Technical speaking on AI-assisted engineering, incident response, agentic systems, and distributed systems.
- IEEE-related peer review activity documented on Web of Science, hackathon judging, technical evaluation, review activity, and startup advisory.
- Regular architecture and design reviews across API contracts, schema governance, and Kafka topic hierarchy.
- Mentoring and onboarding engineers in distributed systems, streaming platforms, and cloud-native architecture.

Skills

Languages

Java, Kotlin, Python

Backend

Spring Boot, REST APIs, Microservices, Event-Driven Architecture, Kafka, Kafka Streams, Debezium, Redpanda

Datastores

Postgres, DynamoDB, Oracle DB, Redis, Solr, MySQL, MongoDB, HDFS, HBase, Hive, ClickHouse

Cloud & Delivery

AWS, EKS, RDS, S3, Lambda, Kinesis, IAM, Kubernetes, Docker, Helm, Terraform, OpenTelemetry, Splunk, Datadog, CloudWatch, New Relic, PagerDuty, CI/CD

Web & Testing

JavaScript, Groovy, Grails, HTML, CSS, JUnit, Mockito, Cucumber, Gherkin

Education

M.S. Computer Science

California State University, Los Angeles

B.E. Computer Engineering

Gujarat Technical University