Gautam Mishra

linkedin.com/in/mishragautam github.com/gautammishra

Professional Summary:

- Senior Software Engineer with 10+ years of experience designing and operating large-scale backend systems, including distributed data pipelines and high-throughput APIs.
- Specialized in building and evolving data-intensive platforms across batch and streaming workloads, with hands-on experience operating systems processing billions of events per day.
- Strong background in backend architecture, data engineering, and production operations, with a focus on correctness, scalability, and long-term maintainability.

Technical Summary:

Programming: Java (17/20, Core +EE), Python; JavaScript/TypeScript (working knowledge)
Frameworks: Spring Boot, Dropwizard, JUnit, Mockito; Flask (Python)
Cloud: IBM Cloud, AWS (EC2, Lambda, IAM, S3, SQS/SNS, ECS, EKS, EMR, AWS Batch)
Data Lake: Airflow, Hudi, Spark
Data Streaming: Kafka
Databases: PostgreSQL, MySQL, ClickHouse
NoSQL: Cassandra, Redis
Monitoring: Grafana, Datadog, OpenSearch, PagerDuty
Front-End: Angular (Internal Tooling)
DevOps: Jenkins, GitHub Actions, Docker, Kubernetes, Terraform, Helm

Professional Experience:

Advanced Software Developer | KForce (Client: Fidelity Investments) | April 2024 - Present

- Designed and implemented nightly Spring Batch jobs to generate firm-wide authorization datasets, improving the reliability and correctness of downstream account-authorization APIs used across multiple systems.
- Drove cross-team alignment by designing and owning a shared authorization library, defining data and API contracts, coordinating rollouts, and resolving integration issues with downstream systems.
- Built and operated highly available REST APIs using Spring Boot on AWS EKS, aligned with Fidelity's cloud architecture, security and compliance standards.
- Delivered the team's first internal UI tool for firm and product data ingestion, eliminating recurring weekly manual uploads and reducing operational overhead for the business.
- Used GitHub Copilot selectively to scaffold Angular components and forms, accelerating UI delivery while retaining full ownership of backend APIs, validation logic, data integrity, and security controls.
- Improved code quality and system reliability by increasing automated test coverage and driving adoption of consistent engineering standards across services.

Senior Backend Developer | IBM (The Weather Company) | Aug 2023 – Mar 2024

- Designed and operated backend systems ingesting multi-terabyte datasets daily and distributing millions of weather and lightning data products to enterprise clients with strict latency, availability, and data-quality requirements.
- Built and evolved data ingestion and processing components supporting hundreds of clients, ensuring correctness and reliability at scale.
- Evaluated architectural alternatives using proofs of concept and quantitative analysis to guide decisions across performance, cost, and operational complexity.
- Implemented WebSocket-based communication in a Java web application, reducing end-to-end latency by ~50% and improving real-time data delivery.

- Designed a multi-region redundancy architecture for lightning data ingestion, increasing system resilience and availability across regions.
- Migrated CI/CD pipelines from Jenkins to GitHub Actions and integrated automated security scanning, reducing build times by ~50% and lowering vulnerability exposure.
- Led the transition from traditional data processing services to a serverless architecture using AWS Lambda, achieving ~25% cost reduction and ~20% throughput improvement.
- Migrated REST APIs from Dropwizard to Spring Boot to improve performance, standardize the backend stack, and simplify dependency management.
- Drove backend design reviews, product readiness, and operational standards, while mentoring engineers on maintainable design and production ownership.

Senior Software Engineer | Klaviyo | Jun 2021 – Jul 2023

- Designed and built large-scale stream processing systems powering customer engagement workflows, processing over 5 billion events per day.
- Contributed to the design, implementation, and optimization of Klaviyo's first data lake built on AWS S3 and Apache Hudi, including ingestion pipelines, data layout decisions, and operational optimizations.
- Designed and implemented a materialized views framework using PySpark, enabling teams to efficiently query large datasets while isolating compute from ingestion workloads.
- Built and operated PySpark job orchestration using Apache Airflow, improving reliability and observability of batch data pipelines.
- Designed and implemented a high-throughput ticketing system generating thousands of unique identifiers per second, supporting event ingestion at scales exceeding 5 billion events per day.
- Created a data repair and replay framework within the data lake and ClickHouse ecosystem, enabling safe correction and recovery for billions of daily events.
- Provided technical mentorship through code reviews, design discussions, and architectural guidance, raising the overall quality and robustness of backend systems.

Software Engineer | IBM | Feb 2017 – Jun 2021

- Led the pivotal migration of TWC's radar and weather data products from on-premise infrastructure to IBM Cloud, acting as a workstream owner for key components.
- Designed and implemented scalable event-driven architecture using AWS SQS, SNS and Lambda, reducing operational costs by 25% while improving scalability.
- Built a highly distributed, low-latency streaming system to modernize decade-old application, ensuring continuity during cloud migration.
- Developed Java Spring Boot microservices deployed on Kubernetes to ingest, process, and distribute weather data across AWS, Kafka, and IBM Cloud Object Storage.
- Built an auto-scalable data ingestion system processing terabytes of data daily, reducing costs by ~50% and improving performance by ~30%.
- Designed and operated real-time REST APIs serving 100,000+ requests per second with response times as low as ~30ms for enterprise customers.
- Championed infrastructure-as-code using Terraform and Helm, improving consistency, repeatability, and deployment reliability.
- Drove adoption of Datadog for fine-grained metrics and alerting, significantly improving production observability and incident response.

Java Developer | Reliance Tech Services | Sep 2013 – Nov 2014

- Implemented the 3D Secure protocol for payment systems, ensuring compliance with Indian payment standards and improving transaction security.

- Engineered an SMS routing and load-balancing service, achieving ~40% higher message delivery efficiency and ~25% lower server load, enabling reliable notifications to ~200M users within a 12-hour window.
- Identified gaps in cross-device usability and introduced responsive design standards, improving consistency and usability across web interfaces used by end customers.
- Designed and implemented a location-based service that generated approximately ₹1M INR in monthly revenue and increased user engagement by ~30% through targeted promotions.

Education:

Master's in Computer Science - Illinois Institute of Technology (2015 - 2016)
Bachelor of Engineering in Computer Engineering - University of Mumbai (2010 - 2013)
Diploma in Computer Technology - Thakur Polytechnic (2007 - 2010)