

Sina Rezaeizadeh

rezaeizadehsina@gmail.com | sinarex.github.io | github.com/SinaRex | linkedin.com/in/sina-rezaeizadeh

Professional Experience

Managing Software Engineer, IBM

July 2025 – Present

- Designed and built a vulnerability remediation orchestration platform using **Prefect**, replacing 12–15 developers with a single-click automated workflow across **400+ production streaming applications**, increasing throughput from 30 to **150 apps/day** and cutting remediation cycles from months to weeks.
- Diagnosed a critical production performance bottleneck by identifying unnecessary deep-copy operations in a **Spring Boot Builder pattern**, reducing time complexity from **$O(n^2)$ to $O(n)$** and significantly improving pipeline throughput in production.
- Built a multi-agent AI system using **AutoGen Studio** to automate production change request checklist generation and validation. The team used it during the POC phase and it caught several gaps that would have caused production issues.
- Leading **Kafka tiered storage enablement** across production topics to offload cold data from brokers. The goal is to reduce broker storage costs by moving cold data to cheaper storage tiers.

Senior Software Engineer, IBM

April 2023 – July 2025

- Architected a **comprehensive error handling platform** using **LogScale/Humio** and custom automation tools to proactively detect, analyze, and remediate failures across distributed systems. This gave the team much better visibility into failures before they escalated.
- Built a **Java agent for Kafka Connectors** using **ByteBuddy**. This added distributed tracing via **Jaeger** to trace where things were slowing down or failing.
- Led **chaos testing** for microservices on OpenShift using **LitmusChaos**. During testing we caught connector failures that were silently failing due to disk space issues on the database.
- Designed a monolithic **Helm Chart** for efficient deployment of **200+ applications** on OCP4.
- Developed a cross-cutting **error handling library** via **AspectJ** within the Kafka Stream client library, standardizing error handling across **350+ KStream applications** so joined messages stop getting silently dropped.
- Built a bash **automation tool** that reduced average app update time from **2 hours to 1 minute** across **200+ apps**. This cut out a lot of manual work for the team.
- Designed and implemented a **React** dashboard for a **Kafka Streaming** test coordinator, cutting regression testing time from **2 weeks to 15 minutes** across 300+ KStream apps.

Software Engineer, IBM

July 2022 – April 2023

- Led development of a **Spring/Angular** web application with an admin dashboard to assist an NGO with relocating refugee families.
- Designed a **Kafka message republishing tool** to capture dropped messages caused by processing errors, built with **Vue.js** and **Spring Boot** with OAuth integration.
- Established a benchmarking environment for the **Redpanda** initiative to compare against Confluent using **Helm** and **Jenkins**.
- Migrated stateful applications from Deployment to **StatefulSet** in **OpenShift** using Helm, which gave us stable storage and proper failover.
- Designed a modular **data management platform** for end-to-end ML workflows using **Kubernetes, Kafka, Jenkins, Cassandra** and other cloud-native technologies.
- Built a scalable **Go** microservice for submitting and monitoring ELT data pipelines on Kubernetes using a **TDD** approach.
- Stress-tested a **Kafka** cluster on Kubernetes, automating configuration testing using Helm.

Skills

Languages: Python, Java, Go, JavaScript, C/C++, C#, Swift, SQL

Technologies: Kafka, Flink, Kubernetes, OpenShift, Helm, Docker, Jenkins, Prefect, AutoGen, React, NextJS, Spring Boot, Vue.js

University of Toronto, Graduated 2021

Bachelor of Science (B.S.) in Computer Science Specialization