

MANSEHAJ SINGH

403-616-2829 | mansehaj.singh@uwaterloo.ca | mansehajsingh.com | [LinkedIn](#) | [GitHub](#)

SKILLS

Languages: Java, Python, Javascript, Typescript, C, C++, SQL, HTML, CSS

Tools/Frameworks: Spring, React, Node, Express, Quarkus, Docker, Kubernetes, JEE, Flask, Ansible, Jest

Technologies: AWS (S3, SQS), PostgreSQL, Cassandra, Elasticsearch, Redis, NiFi, MySQL

EDUCATION

University of Waterloo

Sept 2021 – Apr 2026

BASc in Computer Engineering Honours, 3.9 GPA

Waterloo, ON

- Coursework: Data Structures and Algorithms, Systems Programming and Concurrency, Real Time Operating Systems, Compilers, Databases, Networks

EXPERIENCE

Snowflake

Jan 2025 – Apr 2025

Software Engineer Intern

Bellevue, WA, USA

- Incoming on the **Data Platform** team, expanding Snowflake's data lake initiatives

Constant Contact

Jan 2024 – Apr 2024

Software Engineer Intern - Innovation

Waterloo, ON

- Shipped revamp of contact ingestion pipeline in **Java Spring with Amazon SQS**, responsible for importing **70+ million monthly email and SMS contacts**, resulting in **\$20k monthly revenue increase**
- Lead development of contacts ingestion user interface in **React** to maximize customer conversion rates and track usage characteristics of **200k monthly active users**
- Designed storage scheme in **Amazon S3** for cost effective expiry of stale data using object lifecycle rules to transition files to cheaper infrequent access storage
- Implemented fuzzy matching mechanism to automate correction of incoming contact data by leveraging **Levenshtein Distance** to compare text similarity

Canon

May 2023 – Aug 2023

Software Engineer Intern - Platform and Infrastructure

Waterloo, ON

- Prototyped document allocation scheme in **Elasticsearch** to distribute document load across a cluster of mixed performance servers, **indexing metadata of 7 million medical images per year** for search
- Built **Python** cron service for cleanup and compaction of healthcare imaging metadata in **Cassandra**, leading to a reduction of **30+ million redundant healthcare records**
- Multithreaded Java process to parallelize medical image data transfer over TCP, resulting in a maximum of **16x faster transfer time**
- Containerized services using **Docker** and **Kubernetes** to generate build and deployment machines for new continuous integration pipeline, improving frequency of in-depth defect detection to daily instead of weekly

Canon

Sept 2022 – Dec 2022

Software Engineer Intern - Product

Waterloo, ON

- Created **ETL pipelines in NiFi** to normalize inbound medical image metadata for indexing, maximizing searchability on hundreds of attributes across **2 million medical images per year**
- Improved query speed by leveraging a node-level **write-through cache** to store image metadata, prioritizing most searched fields for caching
- Established system for **distributed locking** across nodes in an application cluster using **Redis** to enforce mutual exclusion for image data partitioned across NFS servers
- Automated mass end to end testing upgrade effort using **Ansible** for deployment automation and **Playwright** to mimic user interaction, **reducing test effort by 3 weeks**

Computer Science Computing Facility

Jan 2022 – Apr 2022

Full Stack Developer Intern

Waterloo, ON

- Spearheaded development of web application responsible for room and desk bookings for **9.2k+ faculty members** in **Python** and **Flask** framework
- Contributed to writing **REST** interface layer for a **database provisioning tool** to abstract **PostgreSQL** database creation, user setup, and granting permissions
- Designed and implemented user interfaces in **React with Material UI** component library to deliver interactive portal for database provisioning tool