

Jonathan Medhanie

johnnymedhanie@gmail.com | 347-677-9860 | [linkedin/jonathanmedhanie](https://www.linkedin.com/in/jonathanmedhanie) | [github/johnnymedhanie](https://github.com/johnnymedhanie)

EDUCATION

University of British Columbia

Vancouver, BC | Graduated April 2021

BACHELOR'S DEGREE, MAJOR MATH AND MINOR COMPUTER SCIENCE

Coursework: Advanced Data Structures and Algorithms; Computer Vision; Computational Optimizations; Databases; Probability; Applied Linear Algebra; Graph Theory

WORK EXPERIENCE

GOOGLE | SOFTWARE ENGINEER

New York, NY | Jan 2022 – Current

- Working on YouTube Search & Discovery team, performed **experiment analysis** to determine minimum detectable effect size for upcoming YouTube features, showing statistical significance while minimizing business cost
- Design and implemented backend infrastructure to support payments for YouTube Courses on YouTube app, collaborated with Google Payments team and Legal to launch in North America and Asia
- Built key transaction, conversion and engagement metrics for YouTube Courses users, storing and efficiently querying **100s GBs** a day, and was able to track changes in user events over time

AMAZON | SOFTWARE ENGINEER

Vancouver, BC | Aug 2020 – May 2021

- Worked on the Alexa Skills team, implemented a new skill promotion feature that would analyze code of **every Alexa skill** for malicious intent, immediately flagging and blocking updates to internal skill repository while notifying security teams
- Led the Node 12 migration effort for our internal **AWS Lambda** runtime environment, which includes code changes for deprecating outdated runtimes and updating internal tools to support new features
- Worked on the Metrics analytics pipeline which aggregated alarm data and generated business reports for stakeholders, providing insight into customer usage during skills development cycle

TWITTER | SOFTWARE ENGINEER INTERN

San Francisco, CA | Jun 2019 - Aug 2019

- Part of the Real Time Storage team, redesigned Manhattan's data exporter pipeline, Twitter's real-time distributed database that stores all crucial data including tweets, user info and media metadata, serving **million of queries per second**
- Optimized Exporter performance by a factor of **200x**, reducing file count from **4.5 million** to **20 thousand** during Export operation as a result of implementing custom MapReduce jobs in Java that splits and merges dataset partitions intelligently
- Tuned Hadoop file system and MapReduce configurations during execution to improve CPU and storage use, reducing total job duration from **4 hours** to **20 minutes** on Twitter's largest cluster

CISCO | SOFTWARE ENGINEER INTERN

Vancouver, BC | May 2018 – Aug 2018

- Part of the Observability team, optimized in-house **AWS EC2** health checker by reducing number requests made by pattern matching IP addresses/port, reducing time to completion from **30 minutes** to **25 seconds**
- Implemented a distributed tracing system for our API endpoint in Golang, allowing observability into transactions, latency and duration of each user's unique request
- Refactored production databases' schema minimizing query duration

SKILLS

Languages: C++, Java, Python, Golang, Julia

Technologies: Hadoop, AWS, GCP, Docker

Databases: SQL, DynamoDB, CassandraDB, RocksDB

Concepts: Distributed Systems, Machine Learning