

ANDREW FAN

☎ 604-441-3666 | ✉ andrewqfan@gmail.com | in [andrewqfan](#)

EXPERIENCE

Databricks

Software Engineer Intern - Incoming

May 2026 – Aug 2026

Vancouver, BC

- Incoming Software Engineer Intern

Amazon Web Services

Software Development Engineer Intern - DynamoDB Storage

Sep 2025 – Apr 2026

Seattle, WA

- **Prevented consensus protocol availability regressions** through architecting real-time leader election timing metrics and automated CI/CD integration that blocks deployments when consensus timing exceeds thresholds
- **Established real-time consensus monitoring infrastructure** that detects lock headbutt frequency and lease wait-out duration, enabling immediate alerting when distributed consensus performance degrades
- **Built comprehensive distributed systems validation coverage** by developing capacity test scenarios measuring availability impact during replica failovers and network partitions

Robinhood Markets

Software Engineer Intern - Storage

May 2025 – Aug 2025

Toronto, ON

- **Slashed database incident triage time by 90%** by architecting the first production Model Context Protocol (MCP) servers with an extensible framework that auto-discovers 120+ tools and enables cross-environment connectivity
- **Expanded schema audit coverage from 3% to 100%** by developing a PR check powered by a custom AI agent that leverages MCP tools to fetch real-time diagnostics to automatically label, annotate, and block risky changes
- **Reduced storage ticket churn by 40%** through building a Slack-native Storage Assistant Bot that provides self-serve access to LLM orchestrators using MCP tools—delivering on-demand database diagnostics in seconds

Enovus Energy

Data Engineer Intern - Reporting & Analytics

May 2024 – Apr 2025

Calgary, AB

- **Facilitated \$300+ million in monthly purchases** by developing a dashboard with Python Streamlit, PostgreSQL for interim data, and Databricks for automated scaling, providing price optimizations and site analytics
- **Processed 500 million rows of data weekly** by designing automated ETL pipelines with Apache Spark to clean, integrate, and normalize diverse formats from Spiral, PI, Romeo, and MarketView to store in Databricks data lakes
- **Saved over 10 hours weekly** by automating crude network optimization with Dijkstra's algorithm, incorporating logic to handle pipeline constraints and volume capacities for enhanced efficiency and accuracy

PROJECTS

[Weave](#) | weaveapp.ca | TypeScript, Next.js, Go

- Conceptualized and built a smart group scheduling app that helps users coordinate availability through an interactive grid interface, simplifying meeting planning and enhancing user collaboration

EDUCATION

University of British Columbia

Bachelor of Science in Computer Science and Business

Expected Completion: May 2027

Vancouver, BC

- **CGPA: 3.9/4.0**, Dean's Honour Roll, Trek Scholar (Top 5%)

TECHNICAL SKILLS

Languages: Python, Go, TypeScript, Java, SQL, HTML, CSS, C++, R

Libraries/Frameworks: Django, Flask, Apache Spark, Pandas, PyTorch, TensorFlow, React

Tools: Bazel, Docker, Kubernetes, Terraform, PostgreSQL, MongoDB, Databricks, Git, GitLab, AWS, Google Cloud