

Micah Wylde

✉ micah@micahw.com | 🌐 www.micahw.com | 📷 mwylde

Education

Wesleyan University

B.A. WITH HONORS IN COMPUTER SCIENCE

Middletown, CT

May 2012

Work Experience

Lyft

SOFTWARE ENGINEER, STREAMING PLATFORM

San Francisco, CA

2018-present

- Developing a real-time streaming platform on Apache Flink, supporting dynamic pricing, ETA, fraud, and other use-cases across Lyft
- Led development on Lyft's Flink Kubernetes Operator, now open-source (<https://github.com/lyft/flinkk8soperator>)
- Contributed to Apache Beam's Flink runner, allowing our developers to write streaming pipelines in Python

Sift Science

TECH LEAD, DATA INFRASTRUCTURE

San Francisco, CA

2016-2018

- Leading a team of 10 software engineers and SREs with a responsibility for building the highly-scalable, reliable, and low-latency infrastructure that powers Sift.
- Helped implement a replica HBase cluster with automated failover mechanism
- Integrated circuit-breaking into HBase, producing a substantial reduction in downtime

TECH LEAD, WORKFLOWS

2016

- Led team of 5 to develop Sift's Workflows product, which allows our non-technical customers to define workflow rules that automatically respond to fraud events, without writing any code. Today most Sift customers rely on Workflows, making hundreds of automated decisions per second.

SOFTWARE ENGINEER

2014-2015

- Developed and launched Sift's Device Fingerprinting product, used today by some of the largest websites on the internet
- Led the backend migration of Sift's web console from a thick Rails app to an API-driven SPA using DropWizard
- Redesigned Sift's Elasticsearch infrastructure to improve scalability and query performance

Quantcast

SOFTWARE ENGINEER

San Francisco, CA

2012-2014

- Built a real-time measurement platform that handles over 200,000 requests per second
- Helped scale a real-time ad bidding platform to 500,000 auctions per second
- Developed features on a high-performance C++ webserver for web measurement and ad targeting

Twilio

ENGINEERING INTERN

San Francisco, CA

Summer 2011

- Built a distributed load testing tool for testing Twilio Client, which provides an API for building VoIP apps, with thousands of concurrent calls to determine the maximum number of calls a single server could support without audio degradation.

Washington University CS Department REU

RESEARCHER

St. Louis, MO

Summer 2010

- Designed and evaluated real-time scheduling algorithms for utility-aware non-preemptable, stochastic task sets using machine learning in C++. Worked under Dr. Chris Gill.

Instructional Media Services, Wesleyan University

PROGRAMMING MANAGER

Middletown, CT

2008-2012

- Maintained classroom multimedia technology and academic computing labs. Programmed and designed AMX-based integrated controllers and touch panels. Implemented a touchscreen-based classroom control system in ruby and javascript. Managed student programmers.

Skills

- **Languages:** Java, Scala, JavaScript, Ruby, Rust, Python, C++
- **Tools:** Flink, Beam, Kubernetes, HBase, Kafka, ElasticSearch, PostgreSQL, Envoy, AWS, MapReduce, Spark
- **Specialities:** Distributed systems, Databases, ML Infrastructure, large-scale data processing

Talks

How Lyft Built a Streaming Data Platform on Kubernetes

STRATA SAN JOSE 2020

Virtual Livestream

March 2020

Running Flink and Beam on Kubernetes

APACHECON NORTH AMERICA 2019

Las Vegas, NV

September 2019

Stream Processing at Lyft

SCALE BY THE BAY 2018

San Francisco, CA

November 2018

Reliable Machine Learning on HBase

APPLIED MACHINE LEARNING @ SCALE MEETUP

Seattle, WA

July 2017

Highly-Available HBase

HBASECON 2017

Mountain View, CA

May 2017

API-driven development at Sift Science

API CRAFT MEETUP

San Francisco, CA

October 2014

Publications

The struggle for safety: effectiveness of caterpillar defenses against bird predation

OIKOS

2015

Scalable Utility Aware Scheduling Heuristics for Real-time Tasks with Stochastic Non-preemptive Execution Intervals

23RD EUROMICRO CONFERENCE ON REAL-TIME SYSTEMS

2011