

Arabelle Hou

software engineer

arabelle310@gmail.com
github.com/arabelle
linkedin.com/in/arabelleh

skills

Apache Kafka • Java + Spring Framework • Python • Linux/Unix
[Chainer](#) • Kubernetes + Helm • Docker • Puppet • Git • JavaScript

experience

Salesforce / Software Engineer

JULY 2018 - PRESENT

[Linux, Java, Python, JavaScript, K8s, Helm, Puppet, Bash]

- Designed a Kafka orchestration and topic management system for customer self-service, reducing our topic creation time from days to minutes.
- Helped build a REST API in Spring Boot that provides visibility to all our Kafka clusters; allowing us to check under-replicated partitions, topics, ACLs, quotas, and configs without having to log in to the specific brokers.
- Integrated the UI layer on top of the API with Salesforce SSO for authn/z.
- Completed a large scale upgrade of over 2000 Kafka brokers and related supporting systems from Kafka 0.11 to Kafka 2.0 with zero downtime.
- Wrote Helm charts to deploy services originally running on bare metal to run in Docker containers in EKS.

Salesforce / Software Engineer Intern

MAY 2017 - AUGUST 2017

[Linux, Java, Python, JavaScript, Bash]

- Templatized Kafka topic deployments to reduce 500+ lines of config.
- Automated creating and deleting Kafka topic access control lists.
- Built a service to periodically query and publish Kafka consumer group lag/offsets to a dashboard so customers can monitor their groups.

Nippon Telegraph and Telephone / Media Intelligence Labs Intern

SEPTEMBER 2016 - DECEMBER 2016

[Python, Chainer]

- Trained deep neural networks and recurrent neural networks for acoustic parameter regression modeling with over 10,000 files of linguistic data to create a Chinese Text-To-Speech system.
- Solved issues with overfitting and processing speed to improve the accuracy of the model and enhance speech naturalness/comprehension

Demonware / Software Engineering Intern

JANUARY 2016 - AUGUST 2016

[Python, C++, Linux, Docker, AWS]

- Worked on server side Kafka applications to continuously deliver data between clients (e.g. gamers) and customers for game statistics.
- Implemented a load-testing framework - with technologies such as Locust, Gevent, and APScheduler - which is capable of simulating 5000+ requests per second to our Kafka cluster.

education

University of British Columbia / B.A.Sc. Computer Engineering

SEPTEMBER 2013 - MAY 2018, VANCOUVER, B.C.