Stream Processing - Job Description

Location: London

Summary:

Yelp connects people with great local businesses.

We are a data-driven company, handling billions of events per day to answer many application and business related questions. At the center of providing scalable solutions to these challenges is Yelp's Stream Processing team. This team is responsible for building systems that join, query and transform petabytes of data streams and provide the computational infrastructure to run large scale streaming jobs across Yelp. All engineering teams at Yelp use the stream processing platform and build application and insights on it.

Yelp engineering culture is driven by our values: we're a cooperative team that values individual authenticity, and encourages "unboring" solutions to problems. New hires are expected to deploy working code their first week -- and your impact will only grow from there with the support of your manager and mentor. At the end of the day, we are all about helping our users, growing as engineers, and having fun in a collaborative environment. We are looking for passionate engineers that enjoy solving challenging problems and are excited about helping us build world-class stream processing technology.

What You Will Do:

- Design and develop Stream Processing infrastructure that will be used by Yelp engineers from across the company.
- Build a large scale distributed, scalable, self service platform that can seamlessly scale to handle billions of events per day.
- Innovate and partner with Product and Data teams to help them build business critical applications and insights on the platform.
- Curate or build reusable blocks to analyze events. Build abstractions and developer tooling that allow engineers to quickly build streaming applications in a self service manner.
- Work on, and contribute to open source software, and have industry impact.

We Are Looking For:

- Strong knowledge of distributed systems and application design, with an understanding of operational and reliability trade-offs.
- Solid foundation in data structures, algorithms and complexity analysis.
- Experience with messaging/queuing systems or stream processing systems - Flink, Kafka, Spark Streaming, Storm, Samza, etc in a high volume environment.
- Experience in creating tools that make developing, monitoring, and debugging software easier.
- Fluency in Python, C++, Java, Scala, Go, or a similar language.
- Ability to collaborate with and influence people at all levels of the organization.
- Excellent written and interpersonal communication skills.

Pluses:

- Experience with configuration management tools (Puppet, Chef, Ansible, Salt, CFEngine), monitoring (Nagios, Sensu, Monit), metrics (Graphite, statsd).
- Experience with AWS.
- Experience working with and contributing to open source projects.

Contact Verena - verena@yelp.com