

# RYAN A. PARK

ryan@ryanpark.org • http://www.ryanpark.org

## EXPERIENCE

**Slack**

**2016-Present**

### Senior Manager, Storage Engineering

I presently lead a team responsible for storage engineering at Slack, including MySQL databases, Solr search clusters, and the Kafka message bus.

**Runscope**

**2013-2016**

### Engineering Manager and Principal Infrastructure Engineer

I was the first employee at Runscope, a debugging and testing service for developers working with APIs. I was the primary engineer responsible for infrastructure and technical operations. As the technical operations team grew, I also began managing our small team of operations engineers.

As a software engineer, I built many of our platform's core microservices. Some of these microservices managed data storage and retrieval, such as a service for parsing our customers' search queries and executing them in Elastic Search. I wrote core platform components including a service framework called Smart Service, a service discovery layer built on ZooKeeper and HAProxy, and a fair job queue backed by Redis. The job queue became Runscope's busiest microservice: it processed 1,000 jobs/sec. and 3,200 HTTP requests/sec., with 99th percentile response times near 75 ms.

I was also responsible for the site's technical operations. I managed operational costs and relationships with Amazon Web Services and other vendors. Since we first launched in 2013, operational costs per customer fell by 71%. I built many tools and microservices to help with monitoring/metrics and our code deployment pipeline. This work relied on open source tools like Puppet, Graphite, and Jenkins, and cloud services like Amazon DynamoDB and Kinesis.

In September 2015 I began managing a small team of operations engineers. I helped set the team's vision and direction, helped with project planning and inter-team communication, and conducted regular one-on-one meetings. Externally, I shared Runscope's tools and learnings with the wider software community, by presenting at conferences like True University and O'Reilly Fluent, and by participating on the AWS Podcast in 2015.

**Pinterest**

**2011-2013**

### Technical Operations Lead

At Pinterest, I was the first Technical Operations hire as it became the fastest-growing social network in history. Pinterest's software was written in Python and used a variety of open-source components like MySQL and Memcached.

I was highly involved in Pinterest's technical operations. I implemented configuration management using Puppet, and I developed a custom REST API to integrate other services with Puppet. I built metric collection infrastructure and monitoring infrastructure. I was also our primary MySQL DBA and migrated all of Pinterest's database servers to solid state disks. This project resulted in vastly improved response times and a 78% cost savings.

At Pinterest, I worked to improve operational efficiency by making better use of technical resources. I personally integrated Amazon Web Services' auto-scaling and spot instance functionality. This was a novel way to use spot instances according to Amazon.com CTO Werner Vogels. I selected a number of infrastructure vendors and negotiated contracts when necessary. My work helped reduce operational cost (per page view) by 64%.

I was also the public face of Pinterest's technical operations. I spoke at a number of conferences including PuppetConf. I was a featured speaker in Dr. Vogels's keynotes at Amazon Web Services re:Invent and the AWS Summit NYC.

# RYAN A. PARK

ryan@ryanpark.org • http://www.ryanpark.org

## **PBworks**

**2009-2011**

### **Head of Operations & IT (2010-2011)**

#### **Front-End Web Engineer (2009-2010)**

As Head of Operations & IT, I was responsible for keeping the PBworks web application online. I was on-call 24/7, and was expected to resolve operational problems quickly and independently. I was responsible for data center hardware operations, Linux administration, MySQL database administration, network management, and procurement. My technical operations projects included optimizing email delivery rates, and reducing disaster recovery time by 2 weeks.

Prior to this, I was a front-end software engineer at PBworks for 18 months. In this role I created and refined features in our customer-facing software product. I spearheaded a usability initiative to simplify onboarding and to improve the conversion rate of free trials. I implemented many features to help customers run large deployments, I resolved scalability issues, and I developed a library for tracking our application's perceived performance. As a software engineer, my primary tools included PHP, HTML5/CSS, and JavaScript.

## **Amazon.com**

**2008-2009**

### **Senior Software Engineer, Retail Customer Experience**

At Amazon, my primary responsibility was to expand the use of AJAX technology on the product page in order to improve the customer experience. I developed new customer-facing features for the detail page, the most highly-visited web page on Amazon.com. These features included a guide to finding replacements for obsolete products, which improved the customer experience and resulted in a measurable increase in sales.

## **Apple, Inc.**

**2007**

### **Summer Intern, Mac OS Performance Engineering**

## **EDUCATION**

**M.S., Computer Science**

**Stanford University**

**B.B.A., Information Systems**

**The George Washington University**