

# ANDREW EHRLICH

1489 Webster St Apt 1307, San Francisco, CA 94115

+1 518-391-3777 ◊ andrew@aehrlich.com

## PHILOSOPHIES

---

Unix Philosophy, Functional Programming, Free Software.

## COMPETENCIES

---

Scala (Cats, ZIO, Circe), Python, Java, Kotlin, Android, Bash, Fish, Linux, Systems Administration, Docker, SBT, Build and Deploy, Git, Functional Programming, Functional Design, Documentation, Diagramming, Data Analysis, SQL, Data Pipelines, Data Warehousing, Change Capture, Hadoop, Kafka, Spark, Systems Analysis, Web Development, Ecommerce.

## WORK EXPERIENCES

---

### **Chartboost, San Francisco CA**

Sep 2019 - Present

*Senior Software Engineer, AdServer Team*

- The AdServer team is responsible for Chartboost's real-time ad serving clusters: coding new features, optimizing ad selection algorithms, reducing response times, and systematically cutting JVM heap usage, using 100% Scala and Google Cloud Platform.

### **Centro, San Francisco CA**

May 2019 - Aug 2019

*Software Architect, Data Team*

- Created third-generation, real-time data platform on Confluent Event Streaming Platform (Kafka), offering sub-second change capture latencies, Avro schema evolutions, and support for a new diverse set of data providers and data consumers across all Centro products.

### **Centro, San Francisco CA**

July 2018 - May 2019

*Principal Software Engineer, Data Team*

- Software Engineering hands-on technical leadership for Scala and Python data pipelines and warehousing services on bare metal hardware running Cloudera's Hadoop platform.
- Scaled system for 3x growth in data volume.

### **Centro, San Francisco CA**

May 2016 - July 2018

*Senior Software Engineer, Data Team*

- Completed Centro's second-generation data architecture serving reports for internal and external customers, using Scala, Spark, Kafka, Impala, Hive, and Cloudera.
- Successful deprecation of Centro's first-generation data architecture and near-complete elimination of customer complaints.

### **Centro, San Francisco CA**

September 2015 - May 2016

*Data Engineer, Data Team*

- Began implementation of Centro's second-generation Kafka-based data ingestion pipelines in Scala.

### **Unified, San Francisco CA**

February 2014 - September 2015

*Platform Engineer, Data Team*

- End-to-end implementation and operation of Unified's ad performance aggregation service, from REST ingestion to central Kafka log, to AWS Redshift data warehouse.

**Kixeye, San Francisco CA**  
*Analytics Engineer, Data Team*

September 2011 - October 2013

- Worked on in-game telemetry ingestion, transformation, and warehousing using Kafka, HortonWorks Hadoop, Hive, and MySQL.

## **EDUCATION**

---

**Cornell University**  
B.S., Applied Economics and Management

May 2008 - May 2011