David **Zaebst**

Software Engineer

(740) 412-7333 | zaebst@gmail.com | www.zaebst.com | zaebst

Education

The City College of New York — *New York, NY*

M.A. in Physics — *Degree Conferred February 2013*

Ohio University — *Athens, OH*

B.S. in Applied Mathematics and B.A. in Physics — *Degree Conferred June 2006*

Skills

Programming	Ruby, Scala, Spark, Java
Databases	Elasticsearch, MongoDB, Redis, Hive, Redshift
Operating Systems	Debian, Ubuntu, OS X, Kali

Work Experience

Run — *New York, NY*

VP Data Engineering — *May 2013 - May 2017*

- Helped grow a small successful startup as the fifth engineering hire.
- Scaled Run's platform for increased volume after the Publicis acquisition.
- Managed ETL used for time series reporting of ad delivery and billing clients.
- Organized the JIRA queue, wrote tickets, and prioritized the items for the Data Engineering team.
- Handled high priority items key to signing new clients and generating revenue. Examples include an API that allowed clients to download their ad serving logs, cost per complete view charging, and viewability reporting.
- Established team standards for testing, build tools, deployment procedures, code reviews, and git workflows.
- Onboarded new developers by setting up accounts in AWS and github.
- Built out amazon infrastructure, setting up emr clusters, ec2 instances, virtual private cloud networks, and security groups.
- Refactored the ETL workflow that ingests json event files and inserts aggregated ad delivery statistics into MongoDB and ElasticSearch. Using Spark's Scala API, the latest implementation can process json files as large as 170Gb.
- Ported budget management code for a high throughput http server from node.js to java with netty. Used threads and concurrent data structures to reduce database write load.
- Created Ruby AASM state machines to run backups and automate the intake and processing of hourly logs.
- Maintained an administration page to track job state, job status, and maintained json endpoints used for alerting when jobs failed.
- Deployed a sharded MongoDB cluster and an ElasticSearch cluster. Each database held more than 2Tb of ad delivery data and delivered on queries from the website with minimal latency.
- Managed backups on the MongoDB and Elasticsearch clusters.
- Upgraded databases and database drivers from Mongo 2.x to Mongo 3.0 and from Elasticsearch 0.9 to Elasticsearch 2.4. Code reviewed the upgrade to Mongo 3.4 drivers.
- Worked on a modern Rails stack including Backbone.js, Bootstrap, Phusion Passenger, Sidekiq, nginx, Redis, and MongoDB.
- Made application endpoints to serve time series data and connected them to the rickshaw javascript library to graph ad delivery over different time intervals.
- Upgraded the portal, the site for managing advertising campaigns, from Rails 3 to Rails 4.
- Added features to the portal including VAST companion ad units, bulk advertisement import, and reach and frequency graphs.
- Designed ad server log ingestion pipeline which took the ad serving records in JSON, transformed them using Scala and scalding and imported the results into a Redshift database optimized for joining advertising user events.
- Developed original click through rate optimization system based on Bayesian statistics using Hive and Ruby.
- Wrote the first implementation of reach and frequency. This calculated the unique people who saw an ad over different time ranges, the average number of times a person would see the ad, and displayed the results in the graphs in the portal.
- Ensured code quality by testing with RSpec, JUnit, ScalaTest, sbt, gradle, and CircleCI.
- Monitored application and server performance with statsd and datadog.
- Installed Sonatype Nexus repository manager to facilitate code reuse.
- Improved security with ssl termination on nginx servers.
- Used standard tools to track down production issues: pry, chrome developer tools, tcpick, screen, nmap.
- Leveraged Ansible to manage user ssh keys, database deployment, VPC setup, and server configurations.

Sumitomo Mitsui Banking Corporation

Jersey City, NJ

Linux Systems Administrator

June 2011-August 2012

- Administered Red Hat and AIX severs used for transaction processing and fraud detection software.
- Coordinated with internal database, networking, and security teams, along with external product vendors to install new applications.
- Developed Perl libraries as CPAN style modules and packaged modules in RPM format. Libraries were designed to lookup the user id specific to an application and environment, initialize the shell environment, and launch processes as daemons.
- Troubleshot issues in a complex environment including Centrify authentication, MQ, MQFTE, Maestro, WebSphere, TSM, ETL, Oracle RAC, and vendor applications.
- Managed creation of new users, groups, and ssh key infrastructure.
- Reduced application start time for the deposit processing server from 27 minutes to 6 minutes by launching processes in parallel.
- Automated application verification procedures that were manually run by operations. The results were automatically sent in an email report to operations and the application specific team.
- Investigated issues with vendor applications written in C, C++, and Java. Liaised with vendors to resolve issues with products.
- Utilized common diagnostic tools such as strace, trace, gdb, pstree, lsof, wireshark, and netstat.
- Patched applications and setup environments for user acceptance tests.

Morgan Stanley

New York, NY

Senior Systems Analyst

August 2009-February 2011

- Maintained and supported seven trading applications with Linux backends and Windows frontends.
- Took ownership of technical and trading issues. Prioritized multiple issues, managed expectations, and ensured resolution.
- Troubleshot applications and investigated complaints to determine their root cause from database slowness, network latency, and application logic flaws.
- Worked with Sybase stored procedures and Perl to generate trader entitlement reports in excel format. Utilized LDAP to check departments against expected entitlements.
- Added workflows to the in house application setup system. The system had a web interface and stored workflows in databases. Each workflow ran a set of functions to automate the installation and configuration of an application.
- Utilized Perl DBI to update trading application configurations and financial product information stored in Sybase.
- Reconciled financial product records by comparing vendor supplied tables to application tables. The program added default configurations to new products and removed old products prior to daily process caching.
- Developed reports in SQL that displayed summaries of trade executions and commission earnings. Setup user groups to receive reports and scheduled their runtimes with autosys.
- Enhanced server and application monitoring using proprietary publisher-subscriber libraries.
- Built a trade re-reporting tool to query published alerts and resubmit trade reports with needed modifications.
- Created command line tools to query multiple trading processes and aggregate order information. Included functions to cancel specific orders.

Bear Stearns / J.P. Morgan

New York, NY

Linux Build Engineer

October 2007-February 2009

- Built and deployed bond calculation libraries on multiple architectures in C and C++ on Linux servers.
- Updated XS modules that connected Perl to C bond pricing libraries.
- Wrote C and C++ test code for libraries using CUnit and CppUnit.
- Installed Cruise Control for continuous integration. Generated configuration XML and Apache Ant files to work with proprietary software build system.
- Designed internal web pages with CGI, AJAX, JavaScript, CSS, and HTML for Apache servers. The pages displayed holiday date changes and provided documentation.
- Created Perl programs to check FTP feeds for vendor-supplied financial calendars, download data, and convert the data to internal format.
- Supported source code control and builds for the fixed income department with proprietary code, CMake, GMake, ClearCase, and SVN.
- Configured Apache and SVN to serve source controlled files internally.
- Developed tools in object oriented Perl to troubleshoot C++ builds.

Interactive Brokers

Greenwich, CT

Systems Analyst

February 2007-October 2007

- Supported mission critical systems used for trading on world markets.