

OBJECTIVE:

Seeking problems of scale in high volume environments.

SKILLS & LANGUAGES:

Python2/3 (bottle, flask, sqlalchemy), Perl5, Rust, Go, Javascript, C, shell,
TDD, Linux, Apache, SQL (Postgresql, MySQL, Redshift), NoSQL (Elastic, MongoDB, Redis, Riak), React,
AWS (EC2, S3, IAM, RDS, SQS, Kinesis, Athena, Glue), Docker, Git, Gerrit, Subversion, Terraform, Chef, Salt, Puppet.

EXPERIENCE:

Lemur Labs, Los Angeles, CA: CEO and Chief Ninjaneer: Feb 2019 – current

Software consultancy for prototyping ideas and testing market fit. Created a custom shim over an online store to drive traffic and track conversions - reverse engineered the online store API, built a python client and flask server to proxy requests with a React+Redux UI.

ZipRecruiter, Los Angeles, CA: Senior Software Engineer: Sep 2015 – Jan 2019

ZipRecruiter connects people with their next great opportunity through job search and automated smart matching.

Improved search times from minutes to sub-200ms for largest and highest paying customers by offloading candidate search from primary RDBMS to Elastic. Improved acquisition, up-sell and retention of large customers with a new resume search feature over 200 million resumes, importing data from non-searchable raw resume data stored in S3. Proposed, specified, built, tested, load tested, multi-tier disaster recovery, documented and deployed with a team of 3. Created new ops path to build machines outside the monolith.

Transferred to data team to manage data ingress and integrity for both products. Reduced primary database contention and improved cache performance by migrating monolith event data out of the primary data store into a parallel path into S3 and data warehouse, updating all upstream producers and downstream consumers incrementally and with zero downtime. Owned table definitions, S3 data layout and tooling to automate metadata for Athena and Spark as we built a data lake around structured S3 data to reduce load on redshift data warehouse.

python, perl, elastic, mysql, redshift, salt, terraform, docker, aws: ec2, S3, sqs, kinesis

AirMedia, Los Angeles, CA: Chief Ninjaneer: July 2014 – July 2015

First Engineering hire building a video distribution and recommendation platform for premium ad supported video. AirMedia provides free distribution in exchange for a post-roll video recommendation unit. Wrote, deployed and maintained the python back-end services and data reporting hosted in AWS with a mongoDB database. Technical lead for investor due diligence. Onboarding, mentoring and code-review.

DemandMedia, Los Angeles, CA: Lead Engineer + SRE + Ops: Aug 2011 – July 2014

Project Incite: 8 months from idea to 2MM product creating live contextual recommendation across eHow.com's 3.2 million pages and members of an exchange service. Utilized rabbitmq, python, thrift, scrapy, aws, redis, riak, postgresql, django, chef, nagios, git, and gerrit.

Won inaugural company Hackathon to capture a million dollar market via targeted pharmaceutical advertising for eHow. Championed and launched product, deployed at scale across eHow.

Implemented server and process improvements for eHow. Improved Pagerank by tracking and analyzing page load times, shaving hundreds of milliseconds off of TTFB to clients and googlebot. Analyzed and upgraded codebase for php 5.3 – 5.5 transition. Trained eHow Engineers on splunk, git, puppet, python.

Open42 / DemandMedia, Los Angeles, CA: Chief Ninjaneer / Lead Engineer: Nov 2010 – Aug 2011

Open42's small engineering team built a high speed, high capacity, cloud based, "big data" data warehouse for analyzing click traffic at Demand Media, geared to provide machine learning algorithms for related content projection and clustering. Demand Media acquired Open42 in Aug 2011 to bring the technology and team in house.

Fourth engineering hire. Served as technical point-of-contact from Open42 to Demand Media. Wrote data analysis jobs utilizing Hadoop. Due diligence of RSSGraffiti code and redesigning as a work queue to improve database locality. Migrated open42 code repository from SVN to git. Migrated the RSSGraffiti SVN repo to git and online bug tracker to local redmine instance. Picked chef as deployment infrastructure and wrote most of the company recipes. Created automated nagios monitoring for all ec2 cloud hosts. Installed and spearheaded adoption of bug tracking (redmine, integrated into git, gerrit and github), code review in gerrit and continuous integration in jenkins. Amazon EC2 AWS configuration and life cycle management. Redis based IP blocking system for legacy product (7MM entries).

The Rubicon Project, Los Angeles, CA: Senior Software Engineer: Oct 2008 - Sept 2010

The Rubicon Project, the world leader in Yield Optimization technology, launched in 2007 with a mission to automate buying and selling across the \$65 billion global online advertising industry. The 'core engineering' team is responsible for all elements of ad scheduling, serving and back end log processing and aggregation.

Rewrote perl ad scheduler to use a custom rules engine instead of SQL bitmasks, speeding up a 10-30 minute process to sub one minute. Monitored and analyzed slow network transfer in scheduling system, improved sync speed by factor of 5. Stabilized log aggregation and

Andrew Grangaard

andrew.grangaard@gmail.com

(818) 464-6712

28 Outrigger St. #7, MDR, CA 90292

pushed to make our data pipeline a valued piece of infrastructure. Migrated monthly aggregation rollup into a streaming calculation after the calculations exceeded the capacity of the database. Reverse-engineered calculations for financial proration then improved them to use exact splitting of fixed length fractional amounts to ensure consistent sums at each level of roll-up, parallelized as first production use of hadoop cluster. Improved layout of SVN repositories, oversaw/assisted with all SVN merges, implemented unit tests for scheduler, added strictures and warnings to scheduler and aggregator code and fixed resulting errors concurrent with mainline development.

Mediaplex, Los Angeles, CA: Software Engineer Manager: Aug 2006 – Sept 2008

Mediaplex, an operational business unit of ValueClick (Conversant), delivers technology solutions to enable online publishers to meet specific business needs. Lead engineer on team of four engineers enhancing a mature online advertising system for moderate-sized publishers. Selected new features, distributed work, designed code, implemented updates, tested and provided live support. Improved accuracy of team's project estimates and delivered releases on time.

Revived MOJO team by providing leadership that empowered engineers to take ownership of their work. Improved communication, cohesiveness and morale by starting weekly one-on-one meetings to track progress, address issues and grow employees. Built solid relationship with systems team through personal communication, follow-through and professionalism - resulting in improved support.

Extended Perl advertisement allocator. Maintained Apache+mod_perl UI and delivery system on RHEL5 Linux environment. Used MySQL database for counts and configuration. Migrated reports and data to an Oracle data warehouse.

COMMUNITY:

Organizer and speaker for LA Perl 2008-current. <http://la.pm.org>

Technical speaker at Social Linux Expo, DevopsDayLA, AllThingsOpen, and multiple YAPC/tPC conferences.

2019 conference presentations: one at ATO and two at tPC2019.

EDUCATION:

California Institute of Technology, Pasadena, CA

Bachelor of Science, Electrical Engineering. DNF