

Andrew A. Raines

28 November 2011

Contact

- Nashville, TN
- 615-807-0005
- me@draines.com
- Usually called “Drew”
- The most recent version of this document available at:
 - <http://www.draines.com/DrewRaines.org>
 - <http://www.draines.com/DrewRaines.html>
 - <http://www.draines.com/DrewRaines.pdf>
 - <http://www.draines.com/DrewRaines.txt>

Objectives

1. Search technology
2. Analytics
3. Automation

Education

Virginia Polytechnic Institute & State University

- Computer Engineering major
- Was class of 1999, but left early for financial reasons with intention to finish later

Skills

- Clojure, Python, Java, Lisp, Perl, zsh, bash, C
- Django, Ring, Compojure, various other web frameworks/libraries.
- ElasticSearch, Solr, Lucene
- Chef, Puppet
- Linux, Mac, Solaris, OpenBSD, FreeBSD, Windows

Work History

Principal Engineer, Sonian (Apr 2010-Present)

- Cloud-based document archiving. Started in 2006-2007. Received \$10M round B funding in early 2011 led by Amazon.
- SAFE Backend Engineer
 - Spent a year on the SAFE archive team. It is a backend which processes imports from various protocols, loads the data to efficient, proprietary, long-term storage based on S3. It then indexes the data to ElasticSearch on EC2/EBS. Application was one of the first large production Clojure codebases.
 - Wrote streaming HTTP upload to S3 using Ring.
 - Wrote a piece of the application which audited import events and stored metrics in Riak. Riak was later removed from the infrastructure so this part of the project never fully materialized.
 - Improved ElasticSearch interaction in the codebase. Implemented bulk indexing. Wrote a utility for allocating indices based on account dataset size.
- Team Lead, Cloud Search
 - Broke off SAFE team to specialize in search technology.
 - Formed team of three (including me) to spearhead migration and stabilization of 300TiB legacy data from failing Ferret-based system to multiple distributed, Lucene-based ElasticSearch clusters. This scheme has allowed us to scale to a petabyte and currently planning for 3x this scale.

- Ongoing work includes shoring up the ES infrastructure with custom Java plugins to meet security requirements, prevent cluster separation during heavy indexing, and refine operational tools to support large numbers of clusters. Also separating indexing from the import system by an API that is institutional controlled to decouple ES upgrades from the rest of the infrastructure.
- Wrote chef recipes for key search infrastructure components. Closely collaborated with operations to fully automate deployments.
- Almost daily experience diagnosing issues with distributed systems at the JVM level.

Software Developer, NotifyMD (Oct 2007-Apr 2010)

- Healthcare telecom firm started in the 1990s to provide devices to physicians that aid call service.
- Rebranded after the turn of the century away from proprietary hardware to a distributed call center model which works with existing phone system.
- I conceived of and led a project to revamp a Coldfusion/SQL Server application that was showing its age. I aggregated data from legacy SQL Server databases to a modern SOLR-based search application. Began the project in Python in 2007 but adopted Clojure shortly after it was released to the public later that year. Allowed me to quickly build an application that pulled archive data in parallel from the RDBMSes and serialize to Clojure s-expressions for fast reindexing.
- Implemented one-touch deployment for all apps in multiple environments.
- Designed and managed local DNS powered by djbdns to support dev operations.

Lead Developer, Cabedge (Apr 2006-Oct 2007)

- Typical agency-style boutique design shop. They hired me in part to transition their small development team to Ruby on Rails. I worked on a few dozen web applications including e-commerce, custom CMS frameworks, and brochures, often discussing requirements directly with

clients. Eventually gave up on Ruby's performance issues and moved us to Django with great success.

Systems Analyst, Vanderbilt University (Sep 2003-Apr 2006)

- Systems Implementation & Development team.
- Lead developer of project to integrate Proofpoint filter software into university's 60,000-mailbox IMAP/sendmail environment.

Software Systems Specialist, Vanderbilt University (Feb 2001-Sep 2003)

- Primary role was a Solaris/Linux sysadmin.
- Maintained a dozen shared Solaris workstations (Ultra 30s/60s) with an E450 serving NFS, email, etc.
- Migrated mail system from sendmail/UW-IMAP to qmail & Courier IMAP.

Technical Support Consultant, Virginia Tech (1996-1999)

- Worked part-time through school.
- Provided primarily network support over the phone to entire campus community. This was on the cusp of Ethernet becoming a standard and before computers shipped with an interface on the motherboard!
- Supported nameless PCs to Macintosh OS 6-9
- Performed occasional housecalls
- Shared a building with the core university datacenter and a gaggle of Internet veterans who helped me cut my teeth on Unix/Linux.

Ancillary Projects

Clojure

- Developer since 2008.
- Had a few patches accepted in 2009.

Genesis Health Technologies

- Current project accepts data from a HUAWEI CDMA module over TCP, decrypts AES-encrypted data, and dispatches to a RESTful HTTP endpoint via OMQ.
- daemontools, ucspi-tcp, Python, C, ZeroMQ.

Pieces of Flair

- One of the original popular Facebook apps.
- Scaled it to 250K daily users, ranking it formerly in Facebook's top 10 apps by popularity.
- Debian GNU/Linux, nginx, Apache httpd, PHP, MySQL, & varnish.
- Acquired by RockYou summer of 2008.

postal

- Library for constructing and sending RFC822-compliant Internet email messages. Wraps the JavaMail package for message and SMTP support.
- Supports sendmail, STARTTLS, and SSL.

esperanto

- ElasticSearch client for Clojure.
- Still experimenting with proper API.

elmer

- A pastebin for the cloud running at <http://p.draines.com>.
- RESTful: Its only dependency for usage is curl.
- Pluggable: Supports filesystem backend by default. Working on streaming S3 support.

clot

- Clojure fork of PircBot.
- Was the original bot for #clojure which was quickly replaced by more eager efforts.

Gmane

- Lars Magne Ingegritسن's public mailing list <-> NNTP gateway.
- I co-moderate with a team of people around the world.

Other Interests

- Running
- Cycling
- Golf
- Photography
- Music