

Summary

Continuing my life-long quest to lower entropy and improve the quality of everything I come in contact with.

If I had a business card, it would say "Does all the things"

Matthew L Daniel

638 Minna ST #3

San Francisco, CA 94103-2773

+1-636-224-8433

mdaniel@gmail.com

PROFESSIONAL EXPERIENCE

SourceClear, Inc

Senior Principal Software Engineer

(Sep, 2015 -- Present) San Francisco, CA

- Updated the Spring Boot command-line client to behave more like a Unix-y tool, switching from interactive-mode only to one driven by command-line flags and a well-written man page. This was the first of many opportunities to drive error-handling as a priority at SourceClear, in order to give the user constructive information instead of cryptic error messages
- Created fully automated CI job for packaging, digitally signing the jar, GPG signing the Debian artifacts, and uploading releases of the command-line client to the download sites. Removed enormous opportunity for typos, allowed one-click releases by anyone with the necessary credentials
- Promoted the use of "git flow" process, and automation around Maven releases to permit `develop`, `master`, `tag`, and hotfixes as formal feature phases
- Migrated the API server away from Mongo toward Postgres JSONB, significantly reducing the deployment complexity and increasing the stability. Created `@PreLoad` and `@PrePersist` JPA hooks to hide the JSONB complexities so it would materialize as a typed object but persist in a managed `@Column`
- Enabled JMX bindings, and created `MBeans` to give greater visibility and control over the system without requiring restarts
- Introduced Sentry error capture to the API and supporting backend services, to help contextualize errors the users were experiencing. Later, Sentry support was added to the web application, and is under consideration for the command-line client
- Migrated auth to use JWT for completely stateless API server, and universal auth for web and CLI
- Identified and resolved two security issues with our API, and conducted forensics to ensure they were not exploited
- Migrated a legacy puppet, docker, EC2 deployment pipeline to Kubernetes, which dramatically lowered the deployment friction, surfaced insight into the behavior of the apps to more teams, and consolidating the number of required EC2 instances to host the entirety of the infrastructure. The move to Kubernetes replaced Ubuntu with CoreOS for its immutability, auto-update, and docker-centric worldview
- Introduced Hashicorp's Vault for automated x509 certificate management and more fine grained secret storage
- Replaced costly Splunk license with Graylog to free the organization from per-gigabyte logging fees, allowing indexing of all application logs (from all regions), plus the `systemd` journal output of the instances, as well as capturing `syslog` output from the Cisco ASA
- Championed replacing GitHub Enterprise and Jenkins with GitLab to bring more robust Kubernetes integration as well as a more declarative CI/CD pipeline
- Created an IntelliJ plugin for Ansible yaml and Jinja2 templates during hack week
- Separate from work, I enjoy [providing Kubernetes support on StackOverflow](#), where (as of June, 2018) I am in the top 20 answerers for the Kubernetes tag
- Separate from work, created a KeePassXC branch that reads the 1Password `opvault` file format for use under Linux

Radius

Lead Software Engineer

(Mar, 2014 -- Sep, 2015) San Francisco, CA

- Reduce production Sentry alerts from many thousands a day to zero in six months
- Refactor code to include execution context in error messages versus "it didn't work"
- Refactor toward less monolithic codebase
- Refactor toward Scrapy instead of in-house framework
- Refactor production deployment toward systematic and transparent builds and deployment activities
- Ease new web data source on-boarding
- Promote Apache Lucene tokenizer/analyzer versus hand-written code
- Coach and troubleshoot technical issues with engineering tools
- Modify AWS IAM policies to permit self-managed access keys

Senior Software Engineer

(Oct, 2013 -- Mar, 2014) San Francisco, CA

- Inform changes to LAN addressing for better integration with AWS VPC, inbound VPN access and better support for Bonjour/AirPlay
- Provision Nexus repository for binary artifact management, with backups to AWS S3
- Debug and strengthen existing Hadoop/Cascading codebase
- Port stand-alone Hadoop jobs to Cascading for better integration with existing code

Guidewire Software, Inc

Senior Software Engineer

(Jun, 2012 -- Oct, 2013) Foster City, CA

- Replaced Spring MVC with JAX-RS to expose REST semantics previously buried under Spring MVC controllers
- Championed use of AngularJS as MVVM framework where possible
- Leveraged Chrome Frame detection and installation script to reduce "wrong browser" support tickets
- Migrated project from Perforce to Git to ease branch creation and facilitate local file state reporting; history was preserved through `git-p4`
- Published multiple documents on the group's Confluence site describing release processes, best practices, getting started with Git, deployment guides and any other helpful bit of knowledge

EDUCATIONAL BACKGROUND

B.S. Computer Science (1994 -- 1999)

Georgia Institute of Technology, Atlanta, GA

RELEVANT SKILLS

Software Methodologies: Agile, Test Driven Development (TDD), UML

Source Control: Git, Mercurial, Perforce, Subversion

Java: Java SE 8, Java EE 6 (EJB, MDB, Servlet, Resource Adapter), Maven 3 (use and plugin creation), Spring 5, Spring Security, JAX-B, JAX-RS, JAX-WS

JavaScript: Angular.js, Casper.js, Jasmine, jQuery, Phantom.js, React, Require.js, Testacular

Python: Celery, Scrapy

Distributed Data: Hadoop / AWS EMR, Kubernetes, Mesos, Spark

Integration: JSON-LD, JSON Schema, RDF, REST, SOAP, WSDL, WS-I, XML, XML Schema (XSD), XML Stylesheet (XSL, XSL-FO, XSLT), XPath, XQuery

Application Server: Glassfish, Jetty, Tomcat

Database: Oracle, PostgreSQL, SQL Server

Operating System: GNU/Linux, mac OS, Windows