Summary

Matthew L Daniel

1722 Dwight Way
Berkeley, CA 94703
+1-636-224-8433
mdaniel@gmail.com

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"

PROFESSIONAL EXPERIENCE

Open Raven, Inc

Principal Software Engineer
(Aug, 2019 -- May, 2024) *San Franciso, CA*

- Provisioned initial AWS resources for the company, including initial GitLab server, its docker registry, and Kubernetes platform upon which it was deployed
- Introduced 1Password for personal and team credential management
- Introduced Sentry error collection and reporting
- Created initial Spring Boot applications and owned their lifecycles from inception out to production
- Created on-premises deployment tooling using client-run docker image and interactive browser tooling to allow the client to install our application without opaque command-line invocations
- Migrated from self-hosted GitLab to GitLab.com along with the docker registry migration
- Integrated GitLab chat-ops into Slack for reduced deployment friction
- Demoed using esbuild to dramatically speed up the React build pipeline, saving substantial CI minutes and lowering developer friction
- Created CloudFormation template to deploy our application upon a self-hosted Kubernetes platform in their own VPC using the AWS CloudFormation Quick Create functionality. Kubernetes and the application deployed in approximately 8 minutes
- Adjusted that CloudFormation template when the company introduced its SaaS offering of that same platform, enabling reuse regardless of who ran the AWS infrastructure
- Created the CloudFormation template to deploy the IAM Roles into the client's AWS account(s) that allowed discovering of their AWS resources
- Contributed to the application for managing the Assume Role connection from our product to the target AWS accounts, either via those deployed IAM Roles or by creating the CloudFormation StackSets at their request
- Authored internal training as well as customer-facing documentation for that same application
- Patched moto to support integration tests against the AWS Organizations API and fixing bugs in it
- Championed the migration away from ElasticSearch to PostgreSQL for ACID behavior and strongly-typed persistence
- Wrote most of the Splunk alerts for identifying problems in both customer clusters as well as our own
- Migrated from self-managed Kubernetes clusters to Elastic Kubernetes Service (EKS)
- Migrated metric collection from Datadog to Amazon Managed Prometheus+Grafana saving the company a great deal of money
- Integrated Firehydrant and Teleport privilege escalation via OpenFaaS for developer self-service during incidents
- Migrated from nginx Ingress Controller to Kong in support of the public API endpoints and authn/authz
- Used Assume Role with Web Identity in GitLab for federated authentication to AWS, reducing static credentials
- Diagnosed and fixed buildx segfaults in GitLab CI when building arm64 docker images
- Laid the groundwork for Google Cloud Platform (GCP) identity federation from our AWS in order to not use static credentials for asset discovery
- Created Google Cloud Shell provisioning mechanism to get close to the "one-click installation" experience
- Laid the groundwork for Azure identity federation from our AWS in order to not use static credentials for asset discovery
- Created Azure Resource Manager (AzureRM) bicep file to get close to the "one-click installation" experience
- Separate from work, I enjoy providing Kubernetes support on StackOverflow, where (as of Nov, 2024) I am in the top 20 answerers for the Kubernetes tag

Atomwise, Inc

Senior Systems Engineer

(Oct, 2018 -- Aug, 2019) *San Franciso, CA*

- Championed the use of PyCharm over text editors for higher code quality
- Migrated git submodules based dependency chain to python wheels
- Migrated AWS-centric tooling toward Kubernetes for better end-user experience
- Used a Kubernetes cluster to replace home-grown multiple VPC isolation and cost accounting system
- Created auto-scaling group for etcd and a separate one for Kubernetes masters to withstand the intermittent load from a 500+ Node spot-instance cluster
- Collected CloudTrail and CloudWatch logs into Graylog for easier search and alerting
- Created Pachyderm machine-learning PoC training system for better reproducibility and visibility
- Used Hashicorp Vault for ssh key leases enabling users to access AWS Batch instances without distributing instance ssh key
- Used Ansible AWX to allow better visibility into playbook runs
- Integrated Atlantis for better visibility into Terraform plan and apply loops

SourceClear, Inc

Senior Principal Software Engineer

(Sep, 2015 -- Oct, 2018) *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 visiblity 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 Intelli/J plugin for Ansible yaml and Jinja2 templates during hack week
- Separate from work, created a KeePassXC branch that reads the 1Password `opvault` file format for use under Linux

EDUCATIONAL BACKGROUND

Bachelor of Science in Computer Science
Georgia Institute of Technology, Atlanta, GA

RELEVANT SKILLS

Software Methodologies: Agile, Test Driven Development (TDD), UML
Source Control: Git, GitHub, GitLab, Mercurial
Java: Java SE 11, Maven 3 (use and plugin creation), Spring 5, Spring Security, JAX-B, JAX-RS, JAX-WS, JPA
Infrastructure: ContainerD / Docker, Kubernetes, AWS, GCP, Azure
Integration: JSON-LD, JSON Schema, REST, XML, XML Schema (XSD), XML Stylesheet (XSL, XSL-FO, XSLT), XPath, XQuery
Application Server: Tomcat
Database: PostgreSQL
Operating System: GNU/Linux, macOS