

Tyler Hampton

SYSTEMS ENGINEER · SRE · PLATFORM ENGINEER · DEVOPS ENGINEER

🏠 [howdoicomputer.lol](#) | 📱 [howdoicomputer](#) | 📺 [howdoicomputer](#)

Summary

Have led growth at infrastructure departments in five different companies. 10+ years of diverse software engineering experience with specialties in large-scale container orchestration systems, polished developer platforms, cloud architectures, and more.

Enthusiastic tinkerer. Perpetual homelabber. Prefers a command line interface environment as a big fan of Emacs, Linux, and open-source.

Work Experience

Grabango

Berkeley, CA, USA

SENIOR SYSTEMS ENGINEER & SITE ARCHITECTURE LEAD

September 26, 2022 - Present

- Mentor and technical lead for DevOps team
- Implemented a HashiCorp Nomad based orchestration system that allocated services across thousands of IoT devices as part of a hybrid infrastructure system bridging physical locations and GCP cloud infrastructure
- Used Dagger and Go to build portable, localized CI/CD pipelines using Docker; used the same Dagger/Go combo to create developer tooling to enable a Platform-as-a-Service model for deploying services to on-premise infrastructure via Nomad
- Lead meetings and facilitated cultural shifts for adopting SRE practices such as SLIs/SLOs and production readiness checklists
- Lead an internal architecture group involving staff and principal engineers to improve the architecture of physical sites and introduce cultural standards for engineering excellence
- Managed GCP Kubernetes clusters via Terraform and enabled a self-service model by building Helm based deployment pipelines using Jenkins
- Heavily refactored SaltStack codebase to improve idempotency and general code quality as well as make use of Salt's more advanced features like orchestration and event-driven actions
- Co-led a migration to Ansible in an effort to abstract application deployment to Nomad and better define the responsibilities of a configuration management system
- Built a monitoring system with Prometheus, Thanos, Grafana, and Consul that used service discovery to auto-configure application metrics collection
- Improved monitoring, resiliency, and performance for complex storage systems built on top of ZFS
- Retooled AI/ML image inference event streaming architecture by replacing Kafka with RedPanda for improved reliability and performance
- Designed extensive caching layers for package types going to physical sites using pull-through Docker registry caches for container images and Nginx proxies for raw package types
- Lead migration from Nexus, internal PyPi, and Google Artifact Registry to Cloudsmith in order to centralize artifact management, improve supported artifact types, reduce complexity, and save costs
- Wrote Golang based CLI tooling for the systems engineering team to automate several day-to-day tasks around ticket management, SaltStack state application, k8s automation, and more
- Designed Terraform usage patterns and code organization and led a multi-month initiative to import hand-spun GCP infrastructure into Terraform state

Shutterfly

Santa Clara, CA, USA

SENIOR SYSTEMS ENGINEER

June 22, 2016 - September 9, 2022

- Technical lead for Cloud Platform team with a head count of 10 engineers
- Onboarded, mentored, and helped grow several engineers and interns while at Shutterfly
- Part of committee of principal, staff, and architect contributors to build a larger architectural model for a migration from data center to AWS
- Migrated over 100 microservices of varying languages to Amazon ECS
- Designed several large, foundational Terraform modules that were used by 200 engineers to deploy Security Groups, Elastic Container Service clusters, IAM entities, VPCs, ALBS, and more
- Designed a AWS global resource testing model that used AWS accounts as test environments for globally available resources like IAM entities
- Wrote a custom reverse proxy in Go to facilitate a maximized inverse cache hit rate in order to support a legacy .NET application
- Wrote custom log ingress filters in Node.js on AWS Lambda for translating log entries into structured log events for Splunk ingress
- Wrote custom tooling around AWS ECS in Python to enforce draining nodes from tasks during abrupt scaling events
- Performed capacity planning in accordance with load testing results and metrics to right size capacity
- Embedded into application teams to assist them with the cloud-native architectural designs of their services
- Designed cloud logging architecture utilizing FluentD, FluentBit, and logging sidecars for log ingestion over PrivateLink into Splunk
- Wrote custom FluentBit plugins in Ruby to support generalized logging event patterns for applications running on ECS
- Wrote Terraform module to abstract AWS Kinesis Streams and then assisted teams with a migration away from Kafka
- Designed immutable OS image pipeline using Jenkins pipelines, Ansible playbooks, and Packer to produce AWS AMIs that fed into EC2 nodes within autoscaling groups
- Wrote custom integrations for running GPU dependent tasks on ECS nodes in order to support complex topographical image generation
- Wrote CLI tooling in Go for Cloud Platform team to automate checking ECS clusters, interacting with internal HTTP endpoints, and more
- Extended and customized aws-azure-login Node.js application to support logging into accounts federated to AWS from Azure AD via SAML
- Co-designed a hub and spoke model that used AWS accounts as a security and management segregation layer for environments with Transit Gateway linking all of the environments together

Opower / Oracle

San Francisco, USA

SOFTWARE ENGINEER, INFRASTRUCTURE AND OPERATIONS

January 12, 2015 - May 23, 2016

- Authored and refactored Puppet manifests to converge data center localized hardware nodes to prepare them as Proxmox hosts
- Co-designed a distributed file system storage array using Ceph to act as shared storage layer for Proxmox VMs
- Designed resilient high-availability MySQL cluster using Heartbeat to float a VIP between multiple masters
- Maintained several applications deployed via Proxmox as VMs that were converged via Puppet
- Wrote a tool in Ruby to migrate employee objects from FreeIPA to Active Directory

InsideVault

San Carlos, CA, USA

DEVOPS ENGINEER

February 24, 2014 - December 15, 2015

- Managed MongoDB cluster used by SEO optimization platform
- Used Jenkins to build testing and deployment pipelines for Scala applications using sbt
- Wrote Chef cookbooks for deploying Apache Mesos cluster members to EC2
- Managed Apache Mesos cluster that was used for distributed batch processing

Apollo Group

San Jose, CA, USA

SYSTEMS ADMINISTRATOR

January 23, 2012 - January 10, 2014

- Wrote a custom deployment web UI using Ruby and jQuery to integrate Chef, AWS EC2 nodes, and Haproxy to serve as the main deployment interface for over 100 engineers for dozens of Java based microservices
- Wrote Chef cookbooks that converged EC2 nodes so that they could run Java based applications
- Rewrote the aforementioned deployment tool as part of a migration from AWS to a data center platform running VMWare VSphere

Education

DeAnza College

Cupertino, CA

COMPUTER SCIENCE AND BUSINESS

Mar. 2010 - Aug. 2012

- Was double majoring in computer science and business

Extracurricular Activity

Code for America

Code for San Jose, San Jose, CA

CORE MEMBER

Jun. 2010 - Jun. 2017

- Mentored and onboarded community engineers and local college students into the Code for America organization
- Project lead and main contributor for CaliDrought - a Node.js API that collected drought information from the U.S. government
- Secondary contributor for CaliDroughtViz - a React application using D3.js to visualize drought impact data for California
- Helped run the chapter through various volunteer activities like managing the treasury and running events

Homelab

A PSEUDO-RELEVANT HOBBY THAT IS LIKE A SECOND JOB

- Three nodes are a k3s cluster running on Raspberry Pis with MetalLB in layer2 mode to serve as a load balancer
- The OS for each k3s node is Ubuntu 20.04 converged with Ansible and the persistent volumes are satisfied by a consumer NAS
- A fourth and primary physical node is a consumer grade hardware box with prosumer hybrid drives in a ZFS mirrored pool. The node is named megamind
- Megamind's underlying operating system is NixOS and immutable deployments to the OS are done via deploy-rs and Nix flakes; failed state applies trigger automatic rollback to last OS state snapshot
- Megamind runs applications and services via Nomad, Consul, Vault
- Individual applications are able to publish their metrics endpoints to Consul and then have their metrics collected via Prometheus and displayed using Grafana
- Cloudflare handles DNS with subdomains routing to specific services via Traefik; Let's Encrypt is used to provide SSL certs via the DNS challenge method
- Homelab serves as a personal PaaS that Elixir and Go services are deployed to

Hobbies

- Board games, rock climbing, science fiction/fantasy books, table top RPGs, programming, 3D printing, guitar, electrical engineering/arduino

Skills

- Bash, Nomad, Terraform, Python, Node.js, TypeScript, Kubernetes, k8s, Elixir, Nix, Docker, AWS, ECS, Helm, TCP/IP, Prometheus, Consul, Vault, Linux, Go, SaltStack, Ansible, Puppet, Chef, git, DNS, GCP, ZFS, Docker, SQL, Kafka, Grafana, cloud, distributed computing, raft, Ruby, Dagger, bpf, JavaScript

Resume Source Code

- <https://git.howdoicomputer.lol/howdoicomputer/resume>