

Objective

A senior software engineering position where I can apply my knowledge, experience, and skills as an engineer to help write quality, high performance applications that solve real-world problems.

Skill Summary

- Thorough experience in backend systems & infrastructure, Python, and Linux.
- Experience with a variety of languages including HTML, CSS, JavaScript, C, C++, various SQL dialects (including MySQL, PostgreSQL), bash shell scripting, and assembly. Some experience with Rust, and interest in learning more.
- Experience with AWS (EC2, S3, RDS, SQS, etc.), GCP (Storage, Compute, Pubsub, networking), Azure.
- Experience with tooling such as git, Docker, Kubernetes, nginx, ELK, etc.

Experience

• Senior Software Engineer at Commure

2019 – 2026

Member of a team responsible for all infrastructure at the company, including Kubernetes clusters, deployment of application code (and development of custom tooling to aid in those deployments), problem-solving during outages and incidents, low-level networking, and more.

Responsible for building out, maintaining, and debugging IPsec / "VPN" tunnels to connect provider networks & data to our systems.

Responsible for custom deployment tooling, maintenance of existing deployments of a healthcare platform. Significant amounts of work with Kubernetes clusters and deployments/software running within them.

This role changed significantly during my time there. The company acquired several other companies, each with products & code of their own. Learned their stacks, including helping debug problems, writing new features, etc. During the industry financial crunch at the end of ZIRP, the team shrunk from a size of about 10 to just me, and later grew to about 10 again. Helped grow new additions by sharing what I knew about the existing systems, bringing several people up to speed.

Responsible for several cost-cutting projects, where optimizing infrastructure / cloud resource usage resulted in >\$10k/month savings.

Responsible for taking an existing audit logging system, which was the largest user (by storage space) of a core database, and growing it: migrated it from a SQL table to Google Cloud Storage (1.6 TiB moved); migrated the code to use a combination of Pub/Sub and GCS to write, and upgraded the format significantly. Modified how some logs were generated that resulted in significant performance gains.

• Senior Software Engineer at HouseCanary

2017 – 2019

Part of a team responsible for back-end data and databases, and building APIs to serve that data up to consumers within the company.

Designed and implemented a Python HTTP API to quickly serve property data based on a wide number of filters to a mobile app, and later a web app.

Worked with and deployed a novel spatial database "Tile38". Also worked with Docker, Python 3 (including asynchronous APIs) and a wide variety of AWS services, including EC2, ELBs, PostgreSQL on RDS, and ECS.

Worked to productionize services developed by the research team, taking them from single-machine proof of concepts to ECS services capable of being easily scaled, deployed, and monitored.

• Backend Software Engineer at Basis Science, Intel

2013 – 2016

Note: Intel acquired Basis Science in March of 2014.

Part of a team of engineers responsible for the server-side code and infrastructure for Basis Science. Implemented several new RESTful JSON API services and endpoints for various features, including mobile clients and internal services and tools. Responsible for devops of services, including deployments, being on-call for production issues, and maintenance of production infrastructure.

Led backend effort to internationalize the server-side code, including implementation of several key features and delegation of work to peers where needed.

Migrated production data to S3 from a mix of older, more complicated storage formats.

- **Software Engineer at Google**

2011 – 2013

Responsible for the design, implementation and testing of several internal services. Implemented web services with HTML, JavaScript and AJAX, using Python with Django and a replicated, failover-capable MySQL database.

Implemented backend services capable of sending and handling RPC calls; implemented code to communicate with several systems managed by other teams.

Responsible for releases of software to production servers, and monitoring and debugging of services maintained by the team, and review of team code.

- **Intern at FactSet**

Summer 2010

Analyzed data in two replicated SQL databases that had drifted out of sync with one another; wrote application to match up and analyze semantic differences in the data; this program's output could aid in decision making to bring databases into sync.

- **Co-op at FM Global**

Spring & Summer 2009

Implemented a graphical C# application to enhance auditor workflow.

Updated automated tests that had become defunct back into a working state to improve automated testing of an internal application.

Implemented a plugin in C# for a CRM application.

Education

- BS in Computer Science at Rensselaer Polytechnic Institute, Troy, NY.
GPA of 3.7 of 4.0 (*magna cum laude*).

May 2011