

CONTACT INFORMATION [lrudolph \(AT\) hmc \(DOT\) edu](mailto:lrudolph@hmc.edu)

EDUCATION **Georgia Institute of Technology**, Atlanta, GA **Jan. 2017 - May 2019**
M.S. Computer Science

Harvey Mudd College, Claremont, CA **Sept. 2012 - May 2016**
B.S. Physics

- Major Concentration in Physics with Computers
- Senior Capstone: *Atomistic Simulations of White Dwarf Dynamics (LLNL)*

SKILLS Go, Python, SQL, Java, bash, git
Terraform, Docker, Kubernetes, Helm, Puppet
Apache Kafka, Apache Flink, Amazon Web Services
Bazel, gRPC, Protobuf, Temporal, Open Policy Agent
Prometheus, Grafana

WORK EXPERIENCE **Principal AI Engineer (VideoAmp)** **Apr. 2026 - present**
I'm a member of the AI Infrastructure team which supports chatbot APIs, Model Context Protocol (MCP) tooling, agentic workflows, evaluation frameworks, and partners with domain teams to make their APIs agent-friendly.

Principal Software Engineer (VideoAmp) **Mar. 2022 - Apr. 2026**
As a member of the Platform Infrastructure team at VideoAmp, my role involved developing core gRPC APIs in Go. I was involved in enhancing our Identity and Access Management systems and resource-sharing APIs. My responsibilities also included modernizing our event-driven architecture (EDA) and establishing developer patterns for projects leveraging EDA. My cross-domain role also involved provisioning new infrastructure using Terraform, as well as writing Open Policy Agent policies in Rego. Additionally, I designed, developed, and maintained Temporal workflows to execute platform business logic and ensure data consistency.

Technologies used: Go, Bazel, Protocol Buffers, gRPC, PostgreSQL, Apache Kafka, Docker, Kubernetes, Helm, AWS, Terraform, Temporal, Open Policy Agent, Datadog, CloudEvents, Snowflake, MCP, Arazzo

Software Engineer (Yelp) **Oct. 2019 - Mar. 2022**
As a member of the Streaming Applications team at Yelp, I worked on maintaining and improving the data streaming infrastructure and interfaces used by Yelp's Kafka-based data pipeline ecosystem, which ingests tens of billions of messages each day. I designed a cost-reduction system for deprecating unused Avro schemas and Kafka topics. I also participated in the on-call rotation as the first-line incident responder, responsible for the full fleet of production Kafka clusters, logging infrastructure, and stream processing infrastructure. Additionally, I contributed to the architecture and development of a custom Kafka Kubernetes operator, as well as the migration of Kafka clusters from EC2 to a k8s-based internal compute platform.

Technologies used: Apache Kafka, Apache Flink, Apache Beam, Python, Go, Java, Scala, Docker, Kubernetes, AWS, Terraform, Puppet, bash, Prometheus, Thanos, Grafana, SignalFx, Splunk

Back-End Developer (BigNerve) **May 2016 - Oct. 2019**
I wrote and maintained code, tests, and documentation for BigNerve's innovator incubator API. I trained new back-end team members and led the development of new API features.

Technologies used: Go, SQL, bash, AWS, Google Cloud Platform, Elasticsearch, Docker

TECHNICAL WRITING **Kafka on Kubernetes at Yelp** **Dec. 2021, Mar. 2022**
I wrote a series of engineering blog posts detailing Yelp's Kubernetes-based Kafka deployment model and migration strategy.
<https://engineeringblog.yelp.com/2021/12/kafka-on-paasta-part-one.html>
<https://engineeringblog.yelp.com/2022/03/kafka-on-paasta-part-two.html>