

CONTACT INFORMATION	lrudolph (AT) hmc (DOT) edu	
EDUCATION	<b>Georgia Institute of Technology</b> , Atlanta, GA <i>M.S. Computer Science</i>	<b>Jan. 2017 - May 2019</b>
	<b>Harvey Mudd College</b> , Claremont, CA <i>B.S. Physics</i>	<b>Sept. 2012 - May 2016</b>
	<ul style="list-style-type: none"><li>• Major Concentration in Physics with Computers</li><li>• Senior Capstone: <i>Atomistic Simulations of White Dwarf Dynamics (LLNL)</i></li></ul>	
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	<b>Principal Software Engineer (VideoAmp)</b> As a member of the Infrastructure & Security Engineering team at VideoAmp, my role involves developing core gRPC APIs in Go. I am involved in enhancing our Identity and Access Management systems and resource-sharing APIs. My responsibilities also include modernizing our event-driven architecture (EDA) and establishing developer patterns for projects leveraging EDA. My cross-domain role also involves provisioning new infrastructure using Terraform, as well as writing Open Policy Agent policies in Rego. Additionally, I design, develop, and maintain Temporal workflows to execute platform business logic and ensure data consistency.	<b>Mar. 2022 - present</b>
	<i>Technologies used:</i> Go, Bazel, Protocol Buffers, gRPC, PostgreSQL, Apache Kafka, Docker, Kubernetes, Helm, AWS, Terraform, Temporal, Open Policy Agent, Datadog, CloudEvents, Snowflake	
	<b>Software Engineer (Yelp)</b> 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.	<b>Oct. 2019 - Mar. 2022</b>
	<i>Technologies used:</i> Apache Kafka, Apache Flink, Apache Beam, Python, Go, Java, Scala, Docker, Kubernetes, AWS, Terraform, Puppet, bash, Prometheus, Thanos, Grafana, SignalFx, Splunk	
	<b>Back-End Developer (BigNerve)</b> 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.	<b>May 2016 - Oct. 2019</b>
	<i>Technologies used:</i> Go, SQL, bash, AWS, Google Cloud Platform, Elasticsearch, Docker	
TECHNICAL WRITING	<b>Kafka on Kubernetes at Yelp</b> I wrote a series of engineering blog posts detailing Yelp's Kubernetes-based Kafka deployment model and migration strategy. <a href="https://engineeringblog.yelp.com/2021/12/kafka-on-paasta-part-one.html">https://engineeringblog.yelp.com/2021/12/kafka-on-paasta-part-one.html</a> <a href="https://engineeringblog.yelp.com/2022/03/kafka-on-paasta-part-two.html">https://engineeringblog.yelp.com/2022/03/kafka-on-paasta-part-two.html</a>	<b>Dec. 2021, Mar. 2022</b>