

JONNY LANGEFELD

SOFTWARE ENGINEER

Bay Area, CA
jonnylangefeld@gmail.com
[GitHub](#)/[LinkedIn](#): jonnylangefeld

I am a data enthusiast and end-to-end developer of complex, Cloud-Native, global software solutions in an environment of backend engineering, robotics & Big Data. My interdisciplinary skill set reaches from software development to dev-ops to data science and machine learning. In my current role at [Cruise](#), I am helping to shape the future by building autonomous driving software, developer productivity tools, infrastructure APIs and automation. My previous hands-on experience in global teams with IBM and Mercedes-Benz is built upon numerous projects to build data-driven products, high throughput APIs and collection, storage, preparation and analysis of data to make valuable predictions and support decisions.

EXPERIENCE

[Cruise](#), San Francisco, CA, USA

Mar. 2020 – Present

Staff Software Engineer, Robotics | Jun. 2023 - Present

- Designed and implemented a distributed tracing system to convey AV task status and progress to human operators. Integrated into Go and Python code; featuring a Protobuf gRPC API.
- Developing software running on the autonomous vehicle to launch the driverless stack with a high degree of automation and resiliency

Staff Software Engineer, Infrastructure | Jul. 2021 - May 2023

- [Successfully transitioned Cruise's central ingress service](#), which handles an average of 40,000 QPS, from self-managed infrastructure to Kubernetes. This migration occurred without any downtime to driverless operations, improved the latency SLO by 65%, and cut costs by 57%.
- Developed and rolled out unified Kubernetes deployment engine. This allowed Cruise to converge multiple systems into one and therefore save on maintenance cost, on-call burden and over 10,000 lines of actively maintained code compared to the previous system.
- Designed and implemented multiple systems to solve [infrastructure resource lifecycle management](#), cluster add-on rollout, alerting setup automation and infrastructure self-service via APIs at Cruise. The system manages over 100 Kubernetes clusters and over 50 helm charts installed depending on configuration and feature flags.
- Saved \$1.6M per year by analyzing data via a Jupyter notebook and identifying dangling infrastructure.
- Mentored Junior, Senior & Staff Engineers, consulted Principal Engineers and conducted over 100 technical interviews.

Senior Software Engineer, Infrastructure | Mar. 2020 - Jun. 2021

- Built infrastructure as code and automation at scale. Automate Cruise's custom cluster creation process, that took engineers initially one week to execute, through a kubernetes controller that does all the work in a matter of one API call.
- Roll out a policy engine across all 40 clusters at Cruise that helps tenants to make their workloads more resilient through enforcing best practices and properties on Kubernetes resources.
- Development of developer productivity tools such as the node pool controller, a kubernetes operator that makes immutable fields on node pool resources mutable through rolling upgrades, saving SREs hours of manual work per cluster during node pool upgrades.

[Mercedes-Benz Research & Development North America](#), Sunnyvale, CA, USA

Apr. 2017 – Feb. 2020

Senior Software Engineer

- Development of Daimler's next generation, high throughput data processing engine, that connects with all Mercedes vehicles and allows GDPR and ISO compliant data exchange between vehicles and backend, consisting of multiple containerized micro services written in Go, running on a Kubernetes cluster
- Development of infrastructure as code with terraform in a hybrid environment of managed resources by cloud providers and self-managed components in a Kubernetes cluster deployed with helm to reduce cloud provider costs by 80%.
- Development of an internal testing framework written in Go, that allows scheduling, execution and reporting of end to end tests from vehicle to backend after processing

Senior Data Scientist

- Development of predictive, classification, clustering models; applying machine learning algorithms on Big Data platforms (with Hadoop, Spark, MLlib, Hive); statistical root cause analysis; data platform project management

[energy-smart](#), Berlin, GERMANY

Nov. 2016 – Oct. 2017

CIO & Co-Founder

- Development of a data platform to serve landlords of larger real estate management firms to track utility usage and deliver utility invoices automatically.
- Responsible for the software development; built a dockerized platform which is able to collect, store and visualize data with a stack of mysql, phpmyadmin, metabase, an sftp server, proprietary software written in Python and a web user interface to manage properties, tenants and connected devices.

[IBM](#), Stuttgart, GERMANY

Oct. 2014 – Mar. 2017

Senior Consultant for Business Analytics | Jan. 2017 - Mar. 2017

Consultant for Business Analytics | Oct. 2017 - Dec. 2016

- Consulting for predictive analytics and business intelligence software (SPSS); development of data-driven software using machine learning algorithms and data mining (pandas, numpy, scikit learn, pyspark, flask, tensorflow).
- Statistical software training for clients; speaker on pre-sales events.

EDUCATION

Bachelor of Science, Business Information Systems

Oct. 2011 – Sep. 2014

[DHBW Stuttgart Cooperative State University](#), Stuttgart, GERMANY

- Co-op study program with IBM consisting of three-month periods of theoretical and practical phases in a continuous exchange. This was the best way to gain a theoretical knowledge foundation and considerable practical business experience at the same time.
- Internships as junior consultant, java developer, strategy consultant and social media campaign researcher

SKILLS

Experience

- Main programming languages are Go and Python; Sufficient in others (Java, Typescript, R)
- Protocol Buffers
- Cloud-Native, Docker, Kubernetes, Helm
- Kubernetes Operators
- Bash, Git, Terraform
- DevOps: Github Actions, Buildkite, Azure DevOps
- Web development (React, Typescript, HTML5, CSS3, PHP)
- Data management (SQL, PostgreSQL, MySQL, MongoDB, phpMyAdmin, IBM DB2, Oracle, Microsoft SQL)
- Predictive analytics (pandas, scipy, IBM SPSS)
- Data science, data mining, machine Learning (CRISP-DM, Chaid, C5.0, Neural Networks, logistic regression, et al.)
- Certified SCRUM Master and Product Owner
- Microsoft Azure, Google Cloud Platform, Mac OS, Linux

Voluntary Activities

- Side projects like [nextux.ai](#)
- [Keynote speaker @Kubecon](#)
- Open-source software ([kubectl-mc](#))
- Video tutorials on Youtube ([Docker](#), [Kubernetes...](#))
- Web administration for a church website
- [Personal blog](#)
- Self-paced learning on udacity (Deep Learning, k8s)
- Developing custom nodes for SPSS Modeler

Qualities

- Entrepreneurial mindset
- Self-directed and passionate
- Leadership and integrity
- Great communicator
- Reasonable decision making