

Chelsea Dole

chelseadole@gmail.com

[Github](#) ▪ [LinkedIn](#) ▪ [Website](#)

TECHNICAL SKILLS

- Python
- Golang
- JavaScript
- SQL
- PostgreSQL
- Apache Kafka
- DynamoDB
- Kubernetes
- Cloud computing (GCP, AWS)
- Docker
- Linux

WORK EXPERIENCE

Brex, Seattle

Staff Software Engineer (Tech Lead)

Aug 2023–Present

Senior Software Engineer

Feb 2022–Aug 2023

Tech lead of the Data Storage team: managing 250+ production PostgreSQL databases, providing organization-level guidance about best practices, and optimizing performance for all teams at Brex.

- Drives long-term direction and architecture for data infrastructure by proactively identifying and addressing problems of scale, reducing cost through hardware/software optimization, writing documentation, and helping engineers to build data-aware systems
- Optimizes read/write performance for PostgreSQL and Redis through query optimization, schema redesign, hash/range partitioning, configuration tuning, building automation, and more
- Contributed to execution of widespread infrastructure cost saving initiatives resulting in 30% and 50% reduction in AWS RDS spend across production and staging respectively
- Reduced a 7TB PostgreSQL database to 800GB through schema redesign and range partitioning with no downtime, resulting in \$38K/year savings and 60% drop in read latency on a critical service
- Pitched, architected, and delivered “DB Metadata Service”, a DynamoDB/Golang-backed gRPC service enabling realtime software/hardware upgrades without impact on metrics or dependencies

Technologies: PostgreSQL, DynamoDB, Redis, Golang, Terraform, PGBouncer, Kubernetes, AWS, Linux

Coffee Meets Bagel, Remote

Senior Data Engineer (Tech Lead)

May 2021–Feb 2022

Data Engineer III

April 2020–May 2021

Data Engineer II

June 2019–April 2020

Architected and built distributed systems to support CMB’s “dater matching” machine learning models and business-critical data warehouses.

- Designed and fully implemented an event-based feature engineering system providing real-time features to machine learning models. Resulted in an 85% reduction in feature calculation time and estimated 75%+ reduction in engineering time needed to author features
- Implemented a company-wide migration from RabbitMQ/Kinesis to Apache Kafka. Pitched project, provisioned infrastructure, and migrated 50+ Kubernetes deployments without downtime or event loss
- Led project to shard a business-critical 3TB Postgres DB with ~4B writes/day using Citus. Built multiple POC sharding options, performed stress testing, and estimated costs

Technologies: Python, PostgreSQL, Kafka, Elasticsearch, AWS, Django, Kubernetes, Redis, Docker, Linux

Zonar Systems, Seattle

Software Development Engineer I

Dec 2017–June 2019

Built backend applications to support a data analysis web platform, including components of a Python data analysis API, a Golang event logging API, and a Dataflow pipeline.

Technologies: Python, Golang, JavaScript/Node.js, Kubernetes, GCP, PostgreSQL, Redis, DataDog, Git

SPEAKING // OPEN SOURCE

[It's Not You, It's Me: Breaking Up with Massive Tables via Partitioning](#) (slides)

Oct 2023

50-minute talk about how to partition live, production-critical PostgreSQL tables while minimizing downtime. Given at: PGConf NYC 2023. Upcoming at: PGConf EU 2023.

[Postgres at Scale: Vacuum, MVCC & Table Bloat](#) (slides)

Mar-Aug 2023

50-minute talk about understanding, quantifying, and mitigating PostgreSQL “table bloat”. Given at: SoCal Linux Expo 2023, PGDay Paris 2023, Nordic PGDay 2023, CitusCon 2023, PASS Data Summit 2023. Upcoming at: Seattle Postgres User Group meetup.

[Understanding Python's Memory Management](#)

Mar 2021

Short introduction to Python memory management. Given at: PyLadies Int'l Women's Day 2021.

[Dreaming of Databases: 5 Questions for Data Layer Decisions](#)

Jan 2020

25-minute talk about database vocabulary and choice. Given at: ThaiPy Bangkok meetup.

[Going Serverless: Deploying with GCP App Engine and Cloud Functions](#)

Jan 2019

30-minute talk about API deployment App Engine. Given at: Puget Sound Python meetup.

PROJECTS // OPEN SOURCE

[Postgres US Foundation](#)

March 2023–Present

Diversity Committee chair for the PgUS Foundation, responsible for distributing \$10K annual funding.

[Invest In Chia](#)

2021–2022

Technical co-founder of a small “green” cryptocurrency startup in the Chia coin space. Built a system which generated and transferred 11PB+ of data to customers across 500+ bare metal servers.

EDUCATION // CERTIFICATIONS

Code Fellows I Seattle, WA

2017

Certificate — Advanced Software Development in Python

University of California, Berkeley

2013–2017

Degree: Political Economy (B.A.)

GPA: 3.40