

Thang Nguyen



EXPERIENCE

Health Rhythms

Senior Software Engineer

New York, NY
Aug. 2019 - Jul. 2024

- Led multiple customer API integrations – adding automatic documentation & contracts to a legacy API using FastAPI, and facilitating project timelines.
- Updated infrastructure to conform to HIPAA compliance standards and remediated security vulnerabilities.
- Developed a Django-based application for vetting hundreds of new prospective users.
- Collaborated with DS to compute individualized hyper parameters used for a proprietary inference.
- Empowered Data Science and API development by defining a schema abstraction layer for DynamoDB.
- Pioneered tool adoption to improve engineering efficiency (e.g. Pydantic, moto, FastAPI, CDK, etc.).
- Profiled and optimized non-performant code (e.g. reduced endpoint execution time by 66%).
- Built CI/CD automation with GitHub Actions (releases, queuing internal job, Docker builds, caching build steps to cut build times by >50 %, optimizing pytest to cut down build time by 20%, etc.).
- Added near realtime system-wide observability and monitoring through Cloudwatch / Datadog integrations.
- Maintained a React dashboard, implementing wireframes, adopting end-to-end browser tests (Cypress), and resolving bugs (e.g. data visualization issues) as needed.
- Accelerated API dev velocity by implementing AWS mocks, parity tests, and Postman collections.
- Established CI, schemas, and tests to a core config service used by mobile and dashboard users.
- Mentored several junior engineers by way of code reviews, pair programming, and debugging.
- Initiated an internal engineering book club discussing long term strategies and goals to the backend engineering group.

The New York Times

Data Engineer

New York, NY
Apr. 2016 - Aug. 2018

- Migrated a streaming data pipeline from AWS to GCP using a Java based message consumer, publishing to Pub/Sub, processing in Dataflow, and persisting to BigQuery for OLAP queries.
- Defined a schema abstraction to manage mappings from raw JSON data to BQ TableRows used in streaming and batch pipelines to simplify schema management and schema evolution requests.
- Ported legacy and new ETL jobs to Airflow, defining best practices along the way.
- Migrated existing sharded tables to partition tables for improved query performance.
- Defined organization level IAM roles to provide granular access to datasets for data governance.
- Implemented a blocking concurrent cache in Java used in a custom Dataflow function.
- Utilized LZ4 compression in a streaming pipeline, while also dropping extraneous data to allow for a better compression ratio (roughly 1/3rd in size, halving monthly costs).

Zipcar

Software Engineer

Boston, MA
Jan. 2013 - Apr. 2016

- Maintained and migrated the legacy Oracle PL/SQL reservation system to a Rails API.
- Utilized the new Rails API and Backbone/FlightJS to create a new point-to-point reservation workflow.
- Optimized legacy SQL queries to improve performance and stability.
- Updated the legacy database image storage and retrieval logic to use a CDN and imagemagick to resize images ahead of time.
- Implemented data pipeline (logs, runtime data) for distributed Rails applications using ELK stack.
- Built various realtime data visualizations using Kibana including: heat maps, dashboards for user behavior, and histograms for profiling database performance.

EDUCATION

Northeastern University

B.S. College of Computer and Information Sciences

Boston, MA
Sept. 2009 - Dec. 2013