

Matt Owen

<http://RAZORBLA.DE>

Phone: +1 (617) 429-4622

Email: matt@worldshadowgovernment.com

LinkedIn: <https://www.linkedin.com/in/whatevergg/>

Mathematical Hacker and [Chaos Engineering](#) Advocate, operating at immense scales with zero downtime. I have worked for 15 years as a Software Engineer and Software Engineering Manager, tackling problems in "Big Data" from the get-go. My teams and I have consistently solved complex problems, owned a ton of infrastructure, and celebrated the delivery of ambitious projects. Whatever my ikigai is, it braids Mathematics and Programming.

tl;dr: Kafka, Airflow, Postgres, Presto, Spark, Kubernetes, Terraform, Python, Golang, *joie de vivre*

Jobs

Owl.co

(New York, New York; January 2023 - September 2023)

Software Engineer

Returning to the startup world, I designed and built products, catering to the Insurance Industry, with a focus on integrating Machine Learning with Human tasks. Here, my role focused on performance of sophisticated models, rather than the high volume, low latency requirements in preceding jobs.

- Led an interdisciplinary team building ML products, automating tasks performed by human investigators: scraping the web, classifying documents, designing ETLs - Clojure, PyTorch, Presto, Spark
- Architected systems to integrate ML inference with human-driven tasks - Clojure, AWS Sagemaker
- Designed and built ETL workflows for analytic databases - Airflow, Spark
- Built dashboards tracking performance of ML models against their human counterparts - Presto, Airflow
- Detected, debugged, and post-mortem'd malfunctioning models
- Built and maintained complex software workflows, orchestrating tasks spanning days
- Owned Agile rituals and meetings: Standup, Planning, Retrospectives, Pre-mortems, and Post-mortems
- Mentored mid-level Engineers on Data Engineering
- Researched and presented data analysis to make go-no-go decisions on new products - Jupyter
- Deployed and maintained infrastructure - AWS CloudFormation, AWS SageMaker

Oracle Data Cloud, MOAT

(New York, New York; December 2019 - October 2021)

Software Engineering Manager

I led a highly technical team to create a vast, event-level data store, used as the source-of-truth for the suite of MOAT products. The real-time system processes 1.2M+ records/second, and requires zero downtime. Consequently, I grew a team with high technical aptitude, and emphasized ownership as a core principle in Software Development.

- Managed and grew team of 7 Data Engineers, ranging from College Recruit to Senior Engineer
- Built multiyear Software Roadmap with Engineering Managers and Product Owners
- Mentored and promoted every Software Engineer on my team
- Collaborated with ML Engineers and Data Scientists to release and update models in production code
- Collaborated with outside Engineering and Data Science stakeholders to design a flexible data pipeline
- Organized and led "Agile" rituals - Sprint Planning, Sprint Review, and Backlog Grooming
- Led project to migrate legacy systems from EC2 to Kubernetes (EKS) - Kubernetes
- Migrated legacy core business logic to modern systems - Kafka, Airflow
- Managed a team owning 800+ instances - AWS
- Managed a budget of \$340,000+ per month
- Authored technical proposals for Data Privacy, System Architecture, and Wire Protocols
- Co-wrote and presented software application proposals, detailing and defending technology decisions
- Reviewed and approved technical design proposals and outage postmortems

Oracle Data Cloud, MOAT (New York, New York; February 2017 - December 2019)

Lead Real-Time Systems Engineer, Senior Data Engineer

I stabilized and scaled a massive computing cluster, halved instance count, and saved over \$2M annually. Comprising a massive 30k codebase, the real-time system contained all business logic to power the MOAT dashboard, and required biweekly deployments. Here, I emphasized stability and correctness, deploying frequent changes across 1,000+ instances.

- Managed weekly software releases for core business logic, contributed to by 4 distinct teams
- Contributed and advised on long-term roadmap as an Independent Contributor
- Onboarded all new hires to MOAT's data pipeline
- Designed and built stream-processing applications processing 1.2M+ events/second - Go, Python, Kafka
- Designed and built system-wide wire protocol - Protobuf
- Built custom software that reduced instance count by 50%, saving over a \$2M dollars - Go
- Built and maintained software end-to-end over 1,000+ AWS instances (c5.xl, r5.8xl)
- Built and maintained real-time processing systems - Go, Python
- Built and maintained Nginx pixel-serving system - Nginx, AWS
- Designed "cold storage" data schema - Parquet
- Maintained historical databases, importing 800,000,000+ rows per day - Highly modified Postgres
- Evaluated replacements for the DB backing the real-time dashboard - Cassandra, Scylla, IgniteDB
- Acquired by Oracle Data Cloud

Chartbeat (New York, New York; December 2014 - December 2016)

Team Lead, Senior Data Engineer

I led an interdisciplinary team as a product-minded Data Engineer, building both the core data pipeline and an initial version of the Chartbeat Historical product. This position introduced me to large-scale distributed systems, leadership, and implementing product-facing changes.

- Led 7-person interdisciplinary Scrum Team
- Organized and led "Agile" rituals - Sprint Planning, Sprint Review, and Backlog Grooming
- Completed Certified Scrum Master Training;
- Designed and built core data pipeline, processing 300,000+ events per seconds - Kafka and Clojure
- Designed and maintained session-level data warehouse - Amazon Redshift
- Designed and maintained sub-second query databases, importing 1,000,000+ rows per hour - Postgres
- Designed wire protocol - Protobuf
- Built and maintained real-time data-scrubbing libraries - Clojure, Java
- Wrote checks, measuring pipeline health and recording instances of data-loss - Nagios
- Deployed and configured production machines - Puppet, Fabric

Harvard University, IQSS (Cambridge, MA; 2009 - 2013)

Lead Developer, Software Maintainer, Statistical Programmer

I maintained the [Zelig Software Suite](#), taught Statistical Programming workshop for Graduate Students at the Institute for Quantitative Social Science, and contributed to the [Dataverse](#). This department, run by Gary King, was at the forefront of "Data Privacy" and "Big Data" before the terms were coined.

- Developed and maintained Zelig Software Suite - R
- Maintained several open-source statistical software packages on CRAN - R
- Built Zelig extension to the Dataverse network - Java
- Taught workshops on Statistical Programming: Beginner, Intermediate, Developing Statistical Packages

Education

Mathematics, focused in Scientific Computing; Bachelors of Science; Boston University; 2009