andrewgrowney.com
github.com/ajgrowney

Andrew Growney

Senior Backend/Platform Software Engineer
Active U.S. Secret Clearance

ajgrowney@gmail.com
913-575-6601
Denver, CO 80215

Professional Experience

Sr Software Engineer - Vantor — October 2021 - Present

- Architect and technical owner of Vantor's Discovery API, the platform powering Vantor's mission-critical consumer facing products, internal research workflows, and derivative geospatial product generation
- Led the evolution of Discovery from a single product ingestion pipeline into a scalable multi-tenant platform securely indexing 100M+ global assets dating back to 1999 with fine-grained authorization control of public, private, and internal datasets
- Designed and operated distributed discovery and ingest infrastructure handling thousands of new assets and real-time event notifications daily across 7 environments with 99.99% production API uptime
- Operated as a cross-functional technical platform lead spanning backend engineering, infrastructure, security remediation, and modernization initiatives
- Reduced our secure asset streaming p95 and p99 latencies by over 85% through architectural improvements to authentication flows, caching strategy, request orchestration, and internal integration collaboration
- Designed and implemented a platform usage telemetry system that has enabled data-driven product improvements, deprecations, and migrations through real-time and historical analysis
- Developed secure service credential management and secret rotation workflows for organization wide CI/CD systems
- Designed and built developer workflows and deployment tooling for low-side development of network-isolated jobs
- Selected to initial AI Champions cohort to help shape internal engineering standards for LLM usage, agent design patterns, and context management strategies

Platform Software Engineer - Atonix Digital — August 2019 - October 2021

- Designed and developed AWS hosted Kubernetes based machine learning engine that handled training, inference, and evaluation of over 90,000 Tensorflow and scikit-learn models for power plant time-series forecasting that saved the company over $2 million annually
- Operated and optimized distributed InfluxDB infrastructure supporting high volume inference retrieval workflows
- Architected and operated multi-environment EKS infrastructure with EFS-backed model storage and spot capacity aware workload orchestration
- Supported end-to-end Kubernetes platform operations including performance testing, cluster upgrades, configuration management, and observability tooling.
- Led company migration from TFVC to Git and redesigned monolithic build system to reusable Azure Pipelines CI/CD templates supporting Angular, .NET, AWS Lambda, and shared library development

Education

University of Kansas - Lawrence, KS — May 2019

- B.S. in Interdisciplinary Computing with Physics concentration and minor in Business Administration
- 3.83 GPA, Honors Program, SELF Engineering Leadership Fellow, Thermography Physics Research Assistant

Languages and Technologies

Languages: Python, Bash, SQL, TypeScript, C++, Groovy, HTML/CSS
Infra: Docker, Linux, Nginx, VPC, Kubernetes, AWS, EKS, Terraform, Helm, Jenkins, Prometheus, Grafana, CDNs, Route53
Backend / Data: PostgreSQL, InfluxDB, Apache Iceberg, Parquet, Protocol Buffers, TensorFlow, scikit-learn, React, REST API design, STAC, GeoJSON, HTTP, TLS, Codex, Augment