

Andrew Caylor

Lakewood, CO | ayyyj@icloud.com | acaylor.dev

linkedin.com/in/andrew-caylor | github.com/acaylor | gitlab.com/acaylor

PROFESSIONAL SUMMARY

Senior Infrastructure, Platform, and Reliability Engineer with 7+ years designing and operating distributed systems across AWS, GCP, Azure, and on-premises environments. Deep experience in Kubernetes, infrastructure as code, GitOps, observability, and internal developer tooling. Operates production EKS at 83 applications, 4,000+ pods, and 100-200 autoscaling nodes.

CORE TECHNOLOGIES

Cloud & Platform: AWS, GCP, Azure, Kubernetes (EKS, GKE, RKE2, OpenShift, AKS), Terraform, Crossplane, Ansible, Helm, KEDA, HPA, VPA

Delivery & Development: ArgoCD, Kustomize, GitHub Actions, GitLab CI, Jenkins, Python, TypeScript, JavaScript, developer CLIs

Reliability & Systems: Prometheus, Grafana, OpenTelemetry, Loki, Mimir, InfluxDB, Alertmanager, Linux, networking, DNS, PKI, TLS, secrets management

Data & Automation: Elasticsearch, PostgreSQL, Kafka, Redis, S3, GCS, self-hosted AI models, AI coding agents, MCP servers

PROFESSIONAL EXPERIENCE

Redox Inc.

Remote

Senior Software Engineer - Platform

February 2022 - Present

- Operate AWS and GCP Kubernetes infrastructure supporting 83 production applications, 4,000+ pods, and an EKS cluster that scales between 100 and 200 nodes.
- Managed successive production EKS upgrades from Kubernetes 1.22 through 1.34 as part of ongoing platform lifecycle ownership.
- Designed and implemented a GitOps continuous-delivery platform using ArgoCD and Crossplane, reducing deployment time by 80%.
- Design application-specific scaling strategies using KEDA, HPA, and VPA according to workload behavior and capacity requirements.
- Lead development of a custom canary deployment and automated rollback workflow for the company's largest application, which operates approximately 1,800 pods.
- Author reusable Terraform modules for AWS and GCP, validating mirrored infrastructure changes in staging before production.
- Built developer-facing CLIs to promote changes into staging, roll back applications during incidents, and manage multiple EKS and GKE environments.
- Led a centralized Prometheus, Grafana, and OpenTelemetry platform with proactive alerting and 20+ multi-cloud dashboards.
- Serve as primary on-call engineer and incident commander, converting production incidents into reliability improvements.
- Authored 15+ architecture proposals and led cross-team initiatives across reliability, delivery, and observability.
- Built LLM-assisted automation for dependency management and platform maintenance, reducing recurring operational toil.

EARLIER EXPERIENCE

Georgia Tech Research Institute

Atlanta, GA

Research Faculty (Cloud Engineer)

September 2021 - February 2022

- Built standardized Kubernetes application platforms supporting research workloads and production services across AWS and Azure.
- Designed and deployed a highly available on-premises Kubernetes cluster integrated with enterprise identity systems and VMware infrastructure.
- Led development of a scalable remote-learning lab platform used by 20+ students per class across three curricula.

Georgia Tech Research Institute

Atlanta, GA

Research Faculty (Systems Administrator)

June 2019 - September 2021

- Automated large-scale bare-metal and VM deployments using Foreman, Red Hat Enterprise Linux, Python, Packer, and Ansible.
- Implemented centralized log aggregation and observability tooling, reducing log-query time from five minutes to under 30 seconds.
- Led migration of critical file storage to a distributed SAN architecture, improving uptime to 99.99%.
- Developed disaster-recovery and business-continuity strategies that reduced MTTR to under one day.

Georgia Tech Research Institute

Atlanta, GA

IT Intern

May 2017 - May 2019

- Managed VMware vSphere storage, networking, and hypervisor infrastructure across three data centers with 1 PB of highly available SAN storage.
- Automated VM configuration, software deployment, and upgrade workflows across Linux and Windows environments.

SELECTED PROJECTS

Multi-Cloud Elasticsearch Platform | **100TB+ indexed data, 2,000+ application clients**

- Led end-to-end design of a Kubernetes-native ECK platform providing a repeatable, cloud-agnostic data service across AWS, GCP, and Azure.
- Authored requirements and architecture documentation, evaluated alternatives, and built prototypes before production rollout.
- Architected a highly available cluster supporting 100TB+ of indexed data and serving more than 2,000 application clients.
- Improved query latency by approximately 20% and throughput by 10% while standardizing monitoring and alerting with Prometheus and Terraform.

High-Scale Metrics and Observability Platform | **42M+ active metric series**

- Re-architected a single-node metrics pipeline as a distributed Kubernetes platform using InfluxDB, Telegraf, and Grafana Mimir.
- Built a Helm-based deployment with S3-backed storage, replacing a custom cloud-init workflow and enabling consistent installations across environments.
- Scaled the platform to 42+ million active metric series in InfluxDB and 12 million series in Grafana Mimir.
- Reduced InfluxDB cold-boot time from approximately 40 minutes to under five minutes and enabled zero-downtime patching for Mimir.

CERTIFICATION

Certified Kubernetes Application Developer (CKAD)

EDUCATION

Kennesaw State University - BBA, Information Systems

May 2019