

Sawyer Czupka

+1 (757) 818-0742 — sawyerczupka@gmail.com — sawyerczupka.com — linkedin.com/in/sawyerczupka

Platform engineer with hands-on experience provisioning and automating cloud and on-premises infrastructure — Kubernetes, Terraform, GitOps, CI/CD — across Azure, AWS, and bare-metal environments, with a strong IaC-first approach from startup to research HPC.

SKILLS

DevOps & Infrastructure

Terraform, Packer, Kubernetes, Docker, Docker Compose, ArgoCD, GitHub Actions, CI/CD, Prometheus, Grafana, Traefik, Nginx, Let's Encrypt, Tailscale, Linux

Cloud, Languages & Tools

Azure (Container Apps, Key Vault, AI Foundry), AWS (Lambda, API Gateway, EC2, S3), Cloudflare, Python, TypeScript, Rust, FastAPI, PostgreSQL, TimescaleDB, Git

EDUCATION

College of William & Mary

Bachelor of Science (B.S.) in Computer Science & Data Science, Cum Laude & Dean's List

Sept. 2021 – May 2025

GPA: 3.6

Relevant Coursework: Cloud Computing, Cybersecurity, Neural Networks, Natural Language Processing

PROFESSIONAL EXPERIENCE

Software Engineering Intern, Luminexis AI / Threat Tec

Aug. 2025 – Dec. 2025

- Served as sole infrastructure owner; provisioned 4-environment Azure stack (Dev, Staging, Demo, Production) using Terraform IaC with shared reusable modules covering Container Apps, PostgreSQL, Key Vault, AI Foundry, and private networking.
- Implemented GitHub Actions CI/CD pipeline automating test, build, and deployment workflows across all environments, eliminating all manual release steps and standardizing delivery.
- Built shared Terraform module library enabling full environment replication in under an hour with complete isolation between instances, reducing new environment provisioning from a multi-day manual process to a single `terraform apply`.

Machine Learning & Software Engineer, Teamculture.ai / L10.tech

Jan. 2024 – Sept. 2024

- Architected serverless AWS infrastructure using Lambda and API Gateway for automated evaluation pipelines; event-driven design with automatic scale-to-zero eliminated idle compute costs while maintaining reliable execution.

Machine Learning Technical Lead, GeoLab @ William & Mary

Jan. 2023 – May 2025

- Operated on-premises Kubernetes deployment on William & Mary HPC cluster; designed microservice topology isolating FastAPI, vLLM inference, and document processing pods for independent resource scaling under shared academic compute constraints.
- Managed full container orchestration lifecycle including pod scheduling, persistent volume provisioning, and service discovery for a multi-component research platform within strict HPC resource quotas.

Machine Learning & Software Engineer Intern, The World Bank & GEF

May 2023 – Sept. 2023

- Deployed document analysis pipeline on W&M HPC shared infrastructure processing 24,000+ project files; gained hands-on experience managing workloads under shared compute resource constraints and academic infrastructure limitations.

PERSONAL PROJECTS

LapEvo: iRacing Telemetry Tracker & AI Racing Coach

Jan. 2025 – Present

- Configured Docker Compose orchestration for local development and Koyeb serverless CaaS for production deployment; implemented GitHub Actions CI/CD pipeline with pytest/testcontainers enabling full integration testing independent of the live iRacing game.
- Architected real-time 60Hz telemetry ingestion pipeline using Rust client with thread-safe async processing and FastAPI/Timescale hypertables; monorepo with auto-generated OpenAPI client (Orval) enforces frontend/backend contract consistency.

Personal Homelab Infrastructure

July 2023 – Present

- Self-host 116-pod Kubernetes cluster across 6 Proxmox VMs with TrueNAS-backed persistent storage via democratic-csi (NFS/iSCSI); all services routed through Traefik reverse proxy with automated Let's Encrypt TLS and Tailscale VPN for secure remote access.
- Fully declarative IaC pipeline: Packer for golden VM image generation, Terraform for VM provisioning, Kubernetes manifests for workload definitions, and ArgoCD for GitOps-based continuous deployment — zero manual steps from commit to running service.
- Manage split-horizon DNS with Technitium (local resolver) and Cloudflare (public DNS), keeping private service records off public infrastructure; Prometheus scrapes all available endpoints with Grafana dashboards for cluster-wide observability.