

(512) 799-4677

DAVE@DGVIGIL.DEV

github.com/dgvigil

linkedin.com/in/dave-vigil

DAVE VIGIL

Platform Engineer | DevOps | Kubernetes | AI Automation

PROFESSIONAL SUMMARY

Platform and DevOps Engineer with 10+ years of experience delivering secure, scalable infrastructure for federal agencies and enterprise organizations. Expertise in Kubernetes (K3s, RKE2, multi-cluster), Infrastructure-as-Code, and DevSecOps – with direct experience supporting classified and unclassified federal environments and eligibility for security clearance based on a DOE background.

Brings a rare combination of federal compliance experience (Platform ONE, 10+ federal entities), open source contribution (rke2-ansible, adopted across federal Kubernetes deployments), and a practical AI automation focus – applying LLMs and workflow tooling to reduce operational toil and accelerate delivery.

TECHNICAL SKILLS

Cloud & Platforms: AWS, Azure, GCP

Infrastructure as Code & Automation: Terraform, OpenTofu, Ansible, Packer, Puppet, Chef

Containers & Orchestration: Docker, Kubernetes (K3s, RKE2, AKS), Helm

CI/CD & Release Engineering: GitHub Actions, GitLab CI, Jenkins, Azure DevOps

Observability & Monitoring: Prometheus, Grafana, custom integrations (Flask/Python)

Programming & Scripting: Python, Bash, Ruby

Practices & Methodologies: GitOps, DevSecOps, Agile

AI & Automation: LLM-assisted development (ChatGPT, Claude, Ollama), workflow automation

CLIENT ENGAGEMENTS

GV5 Technologies LLC (Remote)

Founder/Platform Engineer 2024 – PRESENT

- **Platform Engineering & Automation:** Designing and building automation-first infrastructure and internal platforms using Kubernetes, Terraform/OpenTofu, and modern CI/CD practices to enable repeatable, scalable deployments
- **Application & API Integrations:** Developing applications and services (Flask, Go) that integrate with Slack, Jira, and Google Calendar APIs to automate reporting, generate summaries, and streamline team communication

- **Modular Systems Design:** Building modular, API-driven systems designed for local-first and small-scale AI deployments on commodity hardware, enabling flexible and cost-efficient infrastructure solutions

webAI (Contract, Remote)

DevOps Engineer, Infrastructure & Systems Oct 2024 - Apr 2026

- **Platform Architecture & Kubernetes Strategy:** Designed and maintained multi-cluster Kubernetes platforms supporting distributed AI workloads, establishing scalable patterns for deployment, networking, and cluster operations across teams
- **Infrastructure as Code & Standardization:** Led the development of Terraform/OpenTofu modules and CI/CD workflows that standardized infrastructure provisioning, reducing environment drift and enabling consistent deployments across 5+ engineering teams
- **Cross-Team Enablement:** Partnered with application and platform teams to implement infrastructure best practices, improving developer velocity and reducing onboarding friction for new services
- **CI/CD & Release Engineering:** Built and optimized CI/CD pipelines to support automated build, test, and deployment workflows, improving reliability and accelerating delivery cycles for production services
- **DevSecOps & Reliability:** Integrated security and operational best practices into infrastructure and deployment pipelines, strengthening system reliability and aligning with organizational compliance requirements
- **Operational Ownership & Observability:** Supported day-two operations for Kubernetes environments, contributing to monitoring, troubleshooting, and performance optimization for production AI systems

ServiceNow (Remote)

Systems Engineer FEB 2023 - OCT 2024

- **Process Optimization:** Engineered CI pipeline in GitLab that automated OS regression testing, reducing deployment cycles from 30 days to 3 days—a 90% efficiency improvement saving 270+ hours per release cycle
- **Infrastructure Management:** Created and maintained 15+ Puppet modules managing Dev and Prod environments, implementing repeatable automation solutions that reduced manual intervention by 80%

SUSE Rancher Government Solutions (Remote)

Consulting Kubernetes Engineer / Senior Support Engineer FEB 2021 - FEB 2023

- **Federal Consulting Excellence:** Delivered Rancher, Kubernetes, and CNCF solutions to 10+ federal entities and resolved 200+ technical tickets, facilitating DevSecOps transformation across classified and unclassified environments
- **Open Source Leadership:** Authored RKE2-Ansible configuration management solution (<https://github.com/rancherfederal/rke2-ansible>) and supported full SUSE Rancher product suite including Multi-cluster Manager, Longhorn, and K3s

Sandia National Laboratory / Dept Of Energy (Remote)

Senior Solutions Architect SEPT 2019 - FEB 2021

- **Enterprise Architecture & Innovation:** Led configuration management expansion (Ansible) across national laboratory infrastructure, building CI pipelines enabling 50+ departments to contribute code while designing multi-cluster Kubernetes infrastructure using Rancher
- **Custom Development & Security:** Built custom Molecule driver using Terraform for automated testing, reducing testing overhead by 60%, and developed Python Flask applications for automated monitoring integration across 1000+ endpoints while collaborating with CyberSecurity teams for national security compliance

Rackspace (Remote)

DevOps Engineer JUNE 2018 - SEPT 2019

- **Infrastructure Management:** Built and maintained fully automated CI/CD pipeline for monolithic-to-microservices migration on Google Kubernetes Engine. Managed Infrastructure-as-Code deployments using Helm charts and maintained day-two operations for 500+ Ubuntu servers.

Science Applications International Corporation (Contract, Albuquerque, NM)

Applications Developer / DevOps Engineer MAY 2015 - MAY 2018

- **Application Development & CI/CD:** Coded, deployed, and maintained enterprise Ruby on Rails applications while creating fully automated CI/CD infrastructure for multiple projects using GitLab CI and Jenkins, migrating VM infrastructure to Docker containers
- **Infrastructure Leadership:** Provisioned and maintained 800+ Linux servers and workstations using Puppet, built OpenStack environments, and managed Kubernetes clusters with GPU support for Deep Learning.

AVAILABILITY & ENGAGEMENT

Availability: Available for immediate, long-term W-2 or C2C opportunities.

Security: Eligible for security clearance based on DOE background

Location: Remote-first with occasional travel capability