

Vu Dai An

DevOps Engineer

📞 +84 39 827 0233 | ✉ anvd2401@gmail.com | 🌐 github.com/vudaian24 | 🔗 linkedin.com/in/anvd24
🌐 portfolio.anvd.io.vn | 📍 Ha Noi, Viet Nam

EXPERIENCE

Tokyo Tech Lab Viet Nam

Software Engineer — DevOps & Infrastructure

Oct. 2024 – Present

- Serve as the primary DevOps and infrastructure engineer across three concurrent client products — owning the infrastructure and delivery pipelines for a multi-tenant clinic-booking platform, and supporting the release and environment needs of Kotae (AI chatbot) and StarLive (real-time streaming) in teams of 14–15, working directly with overseas product owners in English.
- Built and maintained delivery pipelines across three CI systems — Jenkins (Jenkinsfile pipelines for container build and publication behind test and quality gates, together with administration of the Jenkins server: agents, credentials, and plugins), GitLab CI, and GitHub Actions.
- Restructured an internal HR platform's GitHub Actions CI into reusable workflow_call components gated on lint and type-check, replacing a workflow_run setup that silently skipped every pull request's first run, and pinned the toolchain with .nvmrc.
- Closed the gap between local, CI, and production environments — replacing MinIO with LocalStack so tests exercise the same S3 API as the deployed service — and enforced commit and pull-request conventions as an automated gate rather than review-time comments.
- Wrote Bash automation for deployment, migration, and environment-provisioning tasks, and served as first-line support for development and release engineers on build and deployment failures, triaging from pipeline and container logs through to root cause.
- Wrote the infrastructure documentation and provisioning runbooks that other engineers follow to onboard new client environments.

Software Intern

Apr. 2024 – Oct. 2024

- Built Gitman, an internal Git repository management tool, from the first commit through to company-wide release — Vue.js and TypeScript interface over a Nest.js and MongoDB API — owning it end to end as the sole developer.
- Shipped Vue 3 and TypeScript workflow features into the production HRM platform used across the company, then moved onto client-facing product work at the end of the internship and was promoted to Software Engineer after six months.

PROJECTS

ERC Booking System – Infrastructure & CI/CD

Jan. 2026 – Jun. 2026

Role: DevOps Engineer / Backend Developer | Tokyo Tech Lab (client product) | Team of 5

Tech: Terraform · AWS (ECS Fargate, Aurora, ElastiCache, CloudFront, WAF, SQS, SES, Backup) · GitHub Actions · Docker · Liquibase

Clinic scheduling platform architected for per-tenant isolation; sole owner of its infrastructure and deployment pipelines.

- Owned the Terraform codebase end to end: 151 resources across 22 reusable modules (VPC, ECS Fargate, Aurora PostgreSQL, ElastiCache, CloudFront + WAF, SQS, SES, Backup, Bastion).
- Designed the per-tenant isolation model so each new clinic gets a dedicated production environment from the same modules, then used it to stand up the first production tenant alongside shared dev/staging.
- Built the GitHub Actions rollout pipeline that patches ECS task definitions, registers revisions, and rolls out API/worker services with stability checks, old-revision cleanup, and path-based selective builds.
- Engineered a branch-gated Liquibase migration pipeline that runs as a one-off ECS Fargate task per environment and auto-generates CloudWatch Logs links for failure triage.
- Authored the S3 + CloudFront release pipeline for the frontend, using immutable cache-control on hashed assets with targeted invalidation for zero-downtime SPA deploys.

Kotae – AI/RAG Chatbot Platform · kotae.ai

Aug. 2024 – Present

Role: DevOps / Full-Stack Engineer | Tokyo Tech Lab (flagship client product) | Team of 14

Tech: Kubernetes · Helm · Docker · AWS Bedrock · OpenAI · Gemini · LangChain · GitHub Actions

Production AI workload serving 10+ SME customers: RAG retrieval over customer knowledge bases, generation backed by AWS Bedrock, OpenAI, and Gemini, operated on Kubernetes across four environments.

- Operated the Helm-managed Kubernetes deployment of this LLM-backed service across four isolated environments (development, test, demo, production), performing release rollouts and rollbacks through Helm against per-environment values, autoscaling, and ingress configuration.
- Kept memory-heavy, bursty AI workloads stable under live traffic: diagnosed pod-level incidents with kubectl — OOMKilled containers, CrashLoopBackOff, failing probes — and tuned resource requests, limits, and probe thresholds.
- Hardened the RAG ingestion path against server-side request forgery, adding private-IP and redirect reachability checks around the crawler, with unit tests covering the blocked ranges.
- Contributed Terraform for the service's Cognito and S3 resources, and unblocked developers when builds and rollouts failed on the shared multi-environment pipeline.

Self-Managed Production Deployment · portfolio.anvd.io.vn

Sep. 2025 – Present

Role: Sole engineer — application, delivery pipeline, and deployment

Tech: Docker · GitHub Actions · GHCR · Nginx · Telegram

Public production service on a self-managed Linux host rather than a managed platform, delivery path owned end to end.

- Containerized the service with a multi-stage Dockerfile — non-root runtime user, container health check — served behind Nginx as a TLS-terminating reverse proxy and pinned to an explicit image version rather than auto-updated.
- Orchestrated CI/CD as chained GitHub Actions workflows on push to main: ci.yml (lint, type-check) gates cd.yml (Buildx build with layer caching, images tagged by commit SHA into private GHCR), which calls deploy.yml — with a manual rollback path redeploying any earlier tag, and outcomes reported to Telegram.
- Deployed over SSH as a least-privilege ci user with the host key pinned in the workflow rather than trusted on first use, recording the outgoing version for rollback and rolling the new one out with compose up --wait behind a retried health check.

SKILLS

Containers & Orchestration: Kubernetes, Helm, kubectl, Docker, Docker Compose

CI/CD: Jenkins, GitLab CI, GitHub Actions, Liquibase

Infrastructure as Code & Scripting: Terraform, Bash, Python, jq

Linux, Virtualization & On-Premises: Linux, systemd, Nginx, Proxmox (VM/LXC)

Cloud (AWS): ECS Fargate, EKS, VPC, Aurora, ElastiCache, CloudFront/WAF, SQS, SES, IAM, Secrets Manager

Monitoring & Observability: CloudWatch Logs, Grafana

Databases: PostgreSQL, MySQL, MongoDB, Redis

Development: TypeScript, Node.js / Nest.js, Next.js, React, Vue 3

PROFESSIONAL SUMMARY

DevOps Engineer with 2+ years delivering and operating containerized workloads across managed cloud and self-managed Linux. Operates the Kubernetes/Helm deployment of a production LLM-backed AI service across four environments and resolves pod-level incidents directly with kubectl. Primary infrastructure owner for a multi-tenant AWS platform — 151 Terraform resources across 22 reusable modules — with zero-downtime container rollouts and gated database migrations, and builds delivery pipelines in Jenkins, GitLab CI, and GitHub Actions. Supports development and release teams through build and deployment failures, and documents the runbooks they follow.

EDUCATION

Dai Nam University

2020 – 2024

Bachelor of Science in Information Technology (GPA: 3.58 / 4.0)

Languages: Vietnamese (native), English (professional working proficiency, used daily at work)