

Dinesh kumar.N

n.dinesh4320@gmail.com | +91 6379327383 | Tiruppur, Tamilnadu | [Github](#) | [LinkedIn](#)

PROFILE

DevOps/MLOps Engineer with around **2+ years of experience** operating **microservice-based and ML-driven applications**. Skilled in orchestrating scalable workloads on **Kubernetes**, implementing automated model lifecycle management **leveraging IaC and CI/CD**. Proven track record on **strategizing deployments, infrastructure scalability, cloud migration and cost optimization**.

PROFESSIONAL EXPERIENCE

DevOps Engineer, Digital Back Office

Apr 2024 - Present

Product - [Musecool](#)

- Migrated a **monolithic 3-tier application** to **4 containerized microservices**, improving system scalability, fault isolation.
- Enforced **auto-scaling for GPU workloads** leveraging **Kserve and Knative** with scale to 0 strategy, reducing infrastructure costs by **~80%** while maintaining on-demand processing performance.
- Integrated **MLflow for experiment tracking and model versioning**, enabling reproducible ML workflows and streamlined promotion of models from experimentation to production.
- Made use of **Redis and RabbitMQ** to buffer and streamline the delivery of music recordings to the inference services preventing overloading.
- Introduced the use of **Optuna for optimizing the model hyperparameters** for the continuously changing input data.
- Structured CI/CD pipelines across **staging and production**, eliminating 90% of manual deployments.
- Strategized cloud migration plan to move infrastructure from **GCP to AWS**, with nearly zero downtime utilizing **Blue/Green** deployment strategy.

Product - [Dataflow](#)

- Architected a **production-grade microservices platform** deployed across **multiple Kubernetes clusters (dev, staging, prod)**, supporting **9+ deployments** with high availability and operational consistency.
- Built and standardized **CI/CD pipelines** for containerized workloads, cutting **manual deployment effort by ~90%** and enabling fast, repeatable releases across AWS and GCP.
- Reduced **Kubernetes pod startup latency by 75%** by implementing container image snapshotting and warm-up instances, significantly improving scale-out responsiveness during traffic spikes.
- Replaced NFS-based Python environment mounts with a **SquashFS-based distribution system** using Kubernetes mount propagation, decreasing **pod initialization time by 60%** and eliminating I/O contention under load.
- Automated infrastructure provisioning using **Terraform and Helm**, enforcing Infrastructure as Code across environments and reliably supporting **100+ concurrent users**.
- Engineered a **periodically scheduled billing pipeline** aggregating data from **6 platform features**, ensuring accurate and reliable periodic billing without manual intervention.

Technologies: Kubernetes, Helm, KServe, Istio, MLFlow, AWS, GCP, Terraform, Git, Bash, Python, Linux

EDUCATION

B.Tech CSE (Cybersecurity)

Amrita Vishwa Vidyapeetham (Deemed University)

Oct 2021 – Apr 2025

SKILLS

- **Cloud Platforms:** AWS, GCP and European Sovereign Cloud providers.
- **Containerization & Orchestration:** Kubernetes, Docker, RabbitMQ, Redis
- **MLOps:** MIFlow, Kserve, Optuna
- **Infrastructure as Code:** Terraform, Ansible, Helm Charts
- **CI/CD & Automation:** GitHub Actions, FluxCD, Bash scripting
- **Programming Languages:** Python, Golang
- **Operating Systems & Networking:** Linux, Istio, Nginx
- **Databases:** Relational databases, ORM concepts
- **Core:** Strategizing Deployments, Cost Estimation and Optimization, Cloud Migration

CERTIFICATIONS

[Certified Kubernetes Administrator](#) (CNCF & Linux Foundation)
[AWS Cloud Practitioner](#) (Amazon Web Services)

Mar 2026 - Mar 2028
Aug 2023 - Feb 2029

PROJECTS

Crack Detection in Tunnels – MLOps Pipeline for Infrared Image Analysis.

- Containerized the trained model and deployed it on **Kubernetes using KServe** for scalable, production-ready inference endpoints.
- Implemented **inference scaling with Knative**, allowing the model service to **scale to zero** during idle periods, **reducing infrastructure costs** while maintaining on-demand crack analysis.
- Implemented **experiment tracking, model versioning, and artifact management using MLflow**, enabling reproducible training and comparison of multiple crack detection models.
- Built an automated **CI/CD pipeline to trigger model build**, validation, and deployment for updated datasets and model versions, improving deployment reliability.
- Enabled **automated reporting workflows** that deliver crack detection results to field engineers for faster inspection and maintenance planning.

Zero Trust Architecture for Microservice-Based Environments

- Designed and implemented a **Zero Trust security model** for **Kubernetes-based microservices**, enforcing **continuous authentication and authorization** using **JWT-based identity validation** and **RBAC-controlled service accounts** across **10+ microservices**.
- **Secured east-west traffic** within the cluster by enforcing **mutual TLS (mTLS)** and **Policy Enforcement Points (PEPs)**, protecting **~99% of inter-service communication** across **multiple namespaces**.
- **Established real-time threat detection and automated alerting**, enabling detection of **policy violations and security breaches** during **10+ test and simulated attack scenarios** in production-like environments.

ACHIEVEMENTS

Cryptoshield Hackathon Tantrotsav'24, Amrita Vishwa Vidyapeetham - Secured 3rd place among 20 teams

NOVA CTF'22, Kumaraguru College of Technology - Secured 5th place among 50 teams

PUBLICATIONS

Zero Trust Security for Web Applications in Microservice-Based Environment - [Link](#)