GANESH SAI DONTINENI

DevOps Engineer | Cloud Infrastructure & Automation
Charlotte, NC (Open to Relocation) | (469) 890-8388 | gsdontineni1998@gmail.com
linkedin.com/in/ganeshsaidontineni | github.com/GaneshSaiD | ganeshsaidontineni.vercel.app

PROFESSIONAL SUMMARY

DevOps and Cloud Infrastructure Engineer specializing in Kubernetes orchestration, Terraform automation, and CI/CD pipeline engineering across AWS, Azure, and GCP. Delivered 30% faster deployments and reduced API latency from 500ms to 150ms at Validus Fintech. Built Kubernetes platforms with GitOps delivery via ArgoCD and Helm, plus an AI-driven self-healing Kubernetes operator with Slack-based human-in-the-loop approval.

EDUCATION

Master of Science, Data Engineering | University of North Texas, TX **Aug 2023 – May 2025**
Post Graduate Student Status Program (PGSSP) | IIIT-Hyderabad **Jul 2020 – Dec 2021**
Bachelor of Technology, Computer Science | JNTU, Hyderabad, India **Jul 2015 – May 2019**

KEY PROJECTS

Kubernetes AI Healer — Self-Healing Kubernetes Operator **2026**
Groq LLaMA 3.3 70B, FastAPI, Prometheus, Kubernetes Python Client, Slack, Streamlit, Docker

- Architected an end-to-end AI-driven self-healing operator for Kubernetes: Prometheus alertmanager webhooks trigger a FastAPI service that invokes Groq LLaMA 3.3 70B for root-cause analysis against live cluster state and event history.
- Designed a Slack-based human-in-the-loop approval gate so no remediation executes without operator sign-off — deliberately trading full autonomy for controlled blast radius in line with production change-management practice.
- Implemented remediation execution through the Kubernetes Python client, covering pod restarts, deployment rollbacks, and resource-limit adjustments, with every action written to a persistent audit trail.
- Built a Streamlit observability dashboard surfacing active alerts, model reasoning, approval status, and remediation history for post-incident review.

Kubernetes GitOps & Observability Platform **2025**
Terraform, Kubernetes, ArgoCD, Helm, Prometheus, Grafana, RBAC, HPA

- Built a reproducible Kubernetes platform on a local cluster using modular Terraform, structuring cluster resources and add-ons as version-controlled, auditable infrastructure code.
- Implemented GitOps continuous delivery with ArgoCD and Helm charts, making Git the single source of truth for cluster state with automated drift detection and reconciliation.
- Configured RBAC policies, namespace isolation, and Horizontal Pod Autoscaling to enforce least-privilege access and demand-based scaling of workloads.
- Deployed a full observability stack with Prometheus and Grafana including custom dashboards, alerting rules, and incident runbooks mapped to each alert condition.

Kubernetes Ops AI Assistant — RAG-Based Documentation Chatbot **2026**
LangChain, OpenAI Embeddings, Pinecone, FastAPI, Streamlit, BeautifulSoup, Docker Compose

- Built a retrieval-augmented generation pipeline over Kubernetes documentation: BeautifulSoup ingestion, LangChain chunking and orchestration, OpenAI embeddings, and Pinecone vector storage for semantic retrieval.
- Served the assistant through a FastAPI backend with a Streamlit interface, containerized with Docker Compose for reproducible local and hosted deployment.
- Reduced time-to-answer for routine Kubernetes troubleshooting questions by grounding LLM responses in retrieved documentation rather than parametric model knowledge.

Heart Disease Prediction at Scale — Big Data ML Pipeline (MS Capstone) **2024**
PySpark, Apache Spark, Random Forest, GCP Dataproc, Pandas, Matplotlib

- Built a distributed ML pipeline using PySpark across the BRFSS dataset of 250,000+ health records, achieving 91% accuracy and a 0.88 weighted F1-score with Random Forest.
- Benchmarked Spark against Hadoop and Cassandra for infrastructure fit, selecting Spark for in-memory processing advantages directly applicable to GCP Dataproc.

PROFESSIONAL EXPERIENCE

Software Engineer — DevOps | Tata Consultancy Services (TCS), Charlotte, NC **Apr 2026 – Present**

- Manage Azure cloud infrastructure and deployment operations for an enterprise banking client, supporting application reliability across multiple environments within an Agile delivery team.
- Maintain and troubleshoot CI/CD pipelines for automated build, test, and deployment workflows, improving pipeline stability and reducing failed builds.
- Develop Python automation for routine operational tasks, eliminating manual effort across release validation and environment configuration.
- Own release coordination and defect resolution, participating in sprint planning, ticket triage, and runbook documentation.

Data Engineer (DevOps) | Cogent Cube Inc., Remote, USA **Jul 2025 – Apr 2026**

- Built and maintained data pipeline infrastructure on AWS, automating provisioning with Terraform and Ansible and integrating CI/CD workflows via GitHub Actions for repeatable, consistent environment builds.
- Containerized data processing workloads with Docker and deployed to Kubernetes across development, staging, and production environments, improving environment consistency and reducing deployment failures.
- Implemented monitoring and alerting with Prometheus and Grafana, building custom dashboards and alert rules that improved pipeline observability and reduced average incident response time.
- Designed ETL workflows ingesting and transforming data from 5+ sources into cloud storage and data warehouses, supporting downstream analytics and reporting.
- Wrote Python and Bash automation for data validation, job scheduling, and routine operational tasks, eliminating roughly 5 hours of manual work per week across environments.

DevSecOps Engineer | Validus Fintech Services, Hyderabad, India **May 2022 – Jul 2023**

- Automated server provisioning and configuration management using Terraform, Ansible, Python, and Bash — increasing deployment rate by 30% and eliminating manual maintenance overhead.
- Resolved CPU and memory bottlenecks and implemented Nginx load balancing — increasing throughput by 35% and reducing API response time from 500ms to 150ms.
- Developed Python scripts for real-time log analysis on AWS EC2 instances, reducing manual monitoring effort by 40% and cutting cloud costs by 20% through proactive issue detection.
- Built GitHub Actions CI/CD pipelines for microservices applications, automating build, test, and deployment stages to enable zero-downtime releases.
- Deployed Prometheus and Grafana alongside custom Datadog dashboards for full-stack observability, reducing MTTR by 25%.
- Configured IAM roles, AWS KMS encryption, and automated security scanning within CI/CD pipelines, reducing vulnerabilities and strengthening compliance posture.

Research Associate / Junior Research Fellow | IIIT-Hyderabad, Hyderabad, India **Jan 2021 – May 2022**

- Built large-scale Python data collection pipelines scraping 15,000+ web pages and social media sources, structuring content for downstream NLP analysis.
- Deployed a domain-specific search engine with query UI, threshold-based classification, and MySQL-indexed document buckets for ranked retrieval.
- Preprocessed and normalized text using BERT and Word2Vec to generate semantic embeddings; mentored junior researchers on data collection and NLP preprocessing workflows.

TECHNICAL SKILLS

Cloud & Infrastructure: AWS (EC2, S3, IAM, KMS, VPC, RDS, API Gateway, CloudWatch), Azure, GCP, Linux Administration, Networking, DNS, TLS

Infrastructure as Code: Terraform, Ansible, CloudFormation, Helm, Modular IaC, Configuration Management

Containers & Orchestration: Kubernetes (EKS), Docker, ArgoCD, GitOps, Nginx, Autoscaling, RBAC

CI/CD & Automation: GitHub Actions, Jenkins, GitLab CI, Git, Blue-Green & Canary Deployments, Python, Bash

Monitoring & Reliability: Prometheus, Grafana, Datadog, ELK Stack, Alerting, Log Aggregation, Incident Response, MTTR Reduction, Runbooks

Data & AI/MLOps: PySpark, Apache Spark, ETL Pipelines, SQL, Scikit-learn, LangChain, RAG Pipelines, OpenAI & Groq APIs, BERT

Languages & Databases: Python, Java, Bash, SQL, C, HTML, CSS | MySQL, PostgreSQL, MongoDB, NoSQL

EARLY EXPERIENCE

Software Development Intern | Rekall Software, Hyderabad, India **Apr 2018 – Jan 2021**

Part-time alongside B.Tech coursework through May 2019; full-time thereafter.

- Automated build workflows and repetitive engineering processes using Python scripts, improving release consistency and reducing manual overhead.
- Developed and optimized backend services using Java and Python across production environments, resolving performance bottlenecks and improving system throughput.