

# GNANESHWAR SREEPATHI

+353 892058601 | [gnaneshwarsreepathi@gmail.com](mailto:gnaneshwarsreepathi@gmail.com) | [LinkedIn](#) | [Medium](#) | [GitHub](#)

## Experience Summary

---

- Accomplished and result-oriented IT professional with **4+ years of experience, in AI-DevOps**. Skilled in supporting, automating, and optimizing mission-critical deployments in AWS, with hands-on expertise in building, deploying, and releasing code across environments. Proficient in source code management, configuration management, and implementing CI/CD pipelines and DevOps best practices to improve system reliability and efficiency.
- Strong advocate of **DevOps and MLOps best practices**, driving efficiency, reproducibility, and reliability in **ML-powered systems**.
- Skilled in building robust **CI/CD pipelines** for **Machine Learning workflows** using **AWS CodePipeline, Jenkins, and GitHub Actions**.
- Experience in Building and maintaining CI pipelines integrated with **ML model registry (MLflow/SageMaker Model Registry)**.
- Implemented multiple **CI/CD** pipelines as a part of **DevOps** role for on-premises and **cloud-based software** using **Jenkins, Ansible, AWS, GCP, Azure DevOps, Docker, and Kubernetes**.
- Experienced in setting up the **Kubernetes cluster on an on-premises data center using Kub spray**
- **Strong Experience with GCP Infrastructure deployments using Terraform**
- Proven expertise in designing and implementing **end-to-end MLOps pipelines** using tools like **MLflow, Kubeflow, SageMaker, and Airflow** for automated model training, testing, and deployment.
- Proficient in **container orchestration** using **Docker, Kubernetes (EKS), and ECS**, enabling scalable and fault-tolerant deployments.
- Experienced in **AWS EC2/VPC/S3/SQS/SNS-based automation through Terraform, Ansible, Python, and Bash Scripts**.
- Experienced in **end-to-end Data Science workflows**, including data ingestion, data preprocessing, feature engineering, model building, evaluation, and deployment.
- Strong expertise in **Python for Data Science**, using libraries such as **NumPy, Pandas, SciPy, Scikit-learn, and Statsmodels** for data manipulation, statistical analysis, and machine learning.
- Hands-on experience in **Machine Learning algorithms** such as Linear & Logistic Regression, Decision Trees, Random Forest, Gradient Boosting, XGBoost, Clustering (K-Means, Hierarchical), and Dimensionality Reduction (PCA).
- Experienced in building and fine-tuning **Generative AI models**, working with **Large Language Models (LLMs)** such as **GPT, LLaMA, Claude, Gemini**, and open-source transformer-based models.
- Expertise in **Prompt Engineering**, prompt optimization, few-shot and chain-of-thought prompting to improve LLM reasoning and response accuracy.
- Experienced using DevOps tools like **Jenkins, GitHub, Ansible, and Docker** to automate and deploy processes to set up and build the environment and applications.
- Involved in **Amazon Cloud Management by setting up Instances, VPCs, and Security Groups, and AWS Security Compliance services, i.e., IAM, Inspector, Key Management Services (KMS), AWS WAF, etc.**
- I have working experience with WebSphere Application Server (WASND), WebLogic, and IBM MQ.
- Experienced with **AWS Services EC2, ECS, ELB, VPC, Cloud Formation, Cloud Front, Cloud Watch, Glacier, IAM, RDS, Route53, S3, SNS, Code Deploying, Kinesis, and Lambda, API Gateway**.
- **Develop** solution options and associate recommendations based on **Technical** and **commercial** considerations and drive the decision-making.
- Developed and optimized complex SQL queries, stored procedures, and user-defined functions to support business logic and reporting
- Experience in working on **IaaS** like **EC2**, **PaaS** like **Docker, RDS**, and **SaaS** like **Cloud Watch**.
- Experience in implementing **CI/CD pipelines from scratch**.
- Extensive experience using **Maven** and **Ansible** as build tools for building **deployable artifacts** from **source code**.
- Hands-on experience in **creating pods** and managing in the **cluster environment** on **Kubernetes**.
- Good working experience on **DevOps tools** such as **Ansible, Maven, SVN, GIT, Jenkins, Docker, and Kubernetes**
- Built ML pipelines using **scikit-learn, TensorFlow, or PyTorch** for classification and regression.
- Built and maintained ETL pipelines using SQL queries in combination with tools like Apache Airflow and AWS Glue
- Automated infrastructure and system administration tasks using **Python scripts** integrated with **AWS CLI, Boto3, and Shell scripts**.
- Good working experience on **DevOps tools** such as **Ansible, Maven, SVN, GIT, Jenkins, Docker, and Kubernetes**
- implemented **LLM-powered applications** using frameworks such as **LangChain, LlamaIndex**, and agent orchestration patterns (ReAct, Tool-based agents, Planner-Executor models).

Experienced in building **Retrieval-Augmented Generation (RAG)** pipelines using vector databases such as **FAISS, Pinecone, Chroma**, enabling contextual and grounded LLM responses.

- Experience in **model evaluation, monitoring, and optimization**, including bias detection, hallucination reduction, and performance tuning for LLM-based systems.
- Knowledge of **MLOps and LLMOps practices**, including model versioning, experiment tracking, deployment strategies, and monitoring using cloud-native tools.
- Expertise in Architecting and Implementing Azure Service Offering, such as Azure cloud services, Azure storage, Azure Active Directory (AD), Azure Resource Manager (ARM), Azure Storage, Azure, Blob Storage, Azure VMs, SQL Database, Azure Functions, Azure Service Fabric, Azure Monitor, and Azure Service Bus.
- Designed and implemented CI/CD pipelines in GitLab CI to automate build, test, and deployment processes across development and production environments.
- Automated deployment of Dockerized microservices using GitLab CI/CD integrated with Helm and Kubernetes

## Core Competencies

---

- Strategic Planning & Execution
- Project Management
- Data Engineering
- Software Development Lifecycle
- Process Improvement
- Business Intelligence

## Technical skills

---

- **Cloud Environments** : AWS and Microsoft AZURE
- **AWS Services** : VPC, Cloud Front, Route 53, API Gateway EC2, RDS, AMI, IAM, S3, Cloud Watch, Cloud Trial, SNS, SQS, Route 53, Cloud Formation, Code Deploy, Elastic Cache, EKS, SFTP, Cognito, ECR, Lambda, Dynamo DB, Auto Scaling, DMS, System Manager, Trusted Adviser, Quick Sight, Data Pipeline, Elastic Search, Certificate Manager, Guard Duty, AWS Cost
- **Azure**: AKS, ACR, Azure AD, virtual Network, virtual machines, storage Azure Account, Load Balancer, key vault, Azure DevOps, Service Bus, Function app, Webapps, Event hub and Notification Hub, AzureSQL, Postgres Cosmos DB, Storage account azure active
- **Machine Learning & MLOps**: Kubeflow, MLflow, SageMaker Model Registry, TensorFlow, PyTorch, Scikit-learn, Airflow.
- **Statistical Methods**: Hypothesis Testing, ANOVA, Principal Component Analysis (PCA), Time Series, Correlation (Chi-square test, covariance), Multivariate Analysis, Bayes Law.
- **Generative AI & LLM**: GPT (OpenAI / Azure OpenAI), LLaMA, Claude, Gemini, Hugging Face Transformers
- **Configuration Management**: Ansible
- **Databases**: MySQL, SQL, PostgreSQL
- **Build Tools**: Maven, Ant, Gradle, MS build, Azure DevOps Build
- **Containerization Tools**: Docker, Kubernetes
- **Version Control Tools**: GIT, GitHub, Subversion
- **Scripting**: Shell Scripting, Python, YAML and Groovy

## Education

---

MSc in data Analytics - National College of Ireland - 2025 – 2026

Ireland

## Experience Summary

---

- **Infosys**  
*Technology analyst* Jun 2021 –Aug 2025.

## Project Summary

---

## Project-1

### Smart Infrastructure Data & AI DevOps Platform

*Client – Larsen & Tubro Construction*

*Role: DS-AI-DevOps*

#### Roles & Responsibilities:

- Designed and deployed highly available, production-grade Kubernetes clusters using Amazon EKS.
- Collaborated with **data engineering and platform teams** to design **AI-ready data pipelines**, using optimized **SQL and Python-based ingestion workflows** to support large-scale model training and inference.
- Implemented **robust data preprocessing pipelines**, including missing value handling, normalization, encoding, and transformation to ensure model stability and consistency.
- Designed and implemented end-to-end **MLOps pipelines** for automated ML model training, testing, deployment, and monitoring using **Kubeflow** and **MLflow**.
- Automated EKS cluster provisioning, node group management, and cluster auto scale configuration using Terraform
- Monitored and troubleshoot workloads in EKS using tools like **CloudWatch, Prometheus, and Grafana**.
- Configured node affinity, taints/tolerations, and pod autoscaling (HPA/VPA) in EKS
- Designed real-time inference pipelines for **generative AI applications**, enabling rapid text generation and interactive user experiences, fine-tuned **large language models (LLMs)** for domain-specific tasks, improving accuracy and relevance in production.
- Built scalable **CI/CD pipelines** for machine learning workflows leveraging Jenkins, GitHub Actions, and GitLab CI.
- Implemented CI/CD pipelines for Terraform using GitHub Actions and Terraform Cloud
- Deployed GPT and other generative AI models on cloud platforms such as **Azure ML**.
- Integrated Terraform with Helm provider to deploy and manage Kubernetes applications
- Established robust **logging and tracking** of ML experiments and **model metrics** using **MLflow**.
- Collaborated with data science teams to operationalize **GPT-based NLP models** and deploy them as REST APIs for production use.
- Led migration of **ML workloads** to cloud-native infrastructure with **Kubernetes operators** and serverless computing.
- Applied **segmentation and clustering techniques** to identify customer cohorts, enabling targeted personalization and AI-driven recommendations.
- Provisioned EKS clusters with **Terraform** and bootstrapped core services (NGINX Ingress, ArgoCD, External-DNS) using Helm via Terraform provider
- Deployed ML models as RESTful APIs leveraging **TensorFlow Serving** and **FastAPI**, facilitating easy access for downstream applications.
- Developed microservices using FastAPI and Flask to serve **AI/ML models** with scalable and low-latency inference.
- Containerized machine learning models using **Docker** and deployed them on **Kubernetes** clusters for high availability and auto-scaling.
- Upgraded and rolled back Helm deployments with zero downtime strategies
- Created and maintained private Helm chart repositories for internal application packaging
- Automated blue/green or canary deployments in EKS with Helm and GitOps workflows using Argo CD

## Project-2

### MLOps Pipeline & AI Model Deployment Platform

*Client – Samsung E&A Arabia Company Limited*

*Role: Sr. DevOps Engineer (MLOps)*

#### Roles & Responsibilities

- Automated the EKS infra deployment using Terraform and GitHub actions Pipeline
- Written the Terraform modules for different services in AWS
- Automated deletion of EBS volumes, snapshots that are older than 365 days using the boto3 Python library
- Worked on deploying the various API primitives in the Kubernetes
- Automated data preprocessing and feature engineering workflows with **Apache Airflow**, ensuring consistent and reproducible **ML pipelines**.
- Configured and maintained infrastructure as code using **Terraform** and **Ansible**, streamlining environment provisioning across cloud and on-premises.
- Worked on developing applications embracing Agile principles
- Worked on the Azure Kubernetes service (AKS) and automated the AKS Infra Deployment using Azure DevOps and

Terraform.

- Written Terraform modules for various services in Azure
- Worked on Cosmos DB deployment using Azure DevOps and Terraform
- Managed **containerized ML workloads** with **Docker** and orchestrated deployments on **Kubernetes**, improving scalability and reliability of ML services.
- Expertise in Architecting and Implementing Azure Service Offering, such as Azure cloud services, Azure storage, Azure Active Directory (AD), Azure Resource Manager (ARM), Azure Storage, Azure, Blob Storage, Azure VMs, SQL Database, Azure Functions, Azure Service Fabric, Azure Monitor, and Azure Service Bus
- Worked with Cloud based hosting solutions (AWS-EC2/S3, RDS, Elastic Beanstalk)
- Worked on the Log Analytics workspace and integrated with the AKS cluster
- Monitored **model performance** and **infrastructure health** using **Prometheus** and **Grafana**, enabling proactive detection of model drift and system failures.
- Migrating on-prem Windows and Linux environments into AWS cloud for better performance and lower usage cost
- Worked on Azure app insights for querying the logs
- Worked on configuring the dashboards
- Implemented AWS Auto scaling groups, Elastic Load Balancing, S3, EBS, EFS, Route 53, VPC, SNS, Cloud Watch, and ec2under AWS
- Used Terraform/CFN scripts to Automate Instances for Manual Instances that were launched
- Created the AWS VPC network for the Installed Instances and configured the Security Groups and Elastic IPs accordingly
- Ensured secure and compliant **ML deployments** by integrating **Vault** for secrets management and enforcing role-based access control (**RBAC**) in **Kubernetes clusters**.
- Utilized EBS to store persistent data and mitigate failure by using snapshots
- Configured and managed security groups, NACL, NAT, and VPC in the Dev/Sandbox environment as per requirements
- Configured AWS cloud infra as code using Terraform and continuous deployment through Jenkins.
- API builds, deployments, and configuration management. Troubleshoot issues related to build and deployment failures - worked on Jenkins and TeamCity
- Implemented **CI/CD workflows** for **machine learning projects** using **GitLab CI** and **Argo CD**, reducing model release cycles
- Wrote Python scripts using Boto3 to automatically spin up the instances in AWS EC2 and OPS Works stacks and integrated with Auto scaling to spin up the servers with configured AMLs automatically.

### Project-3

#### Energy Infrastructure DevOps & Automation Platform

*Client – GAS Vector Saudi Arabia, Ltd*

*Role: Senior DevOps Engineer*

#### Roles & Responsibilities:

- Automated the AKS infra deployment using Terraform and Azure DevOps Pipeline
- Automated GCP Infra deployment using Terraform and GitHub action pipeline
- Automated the infra deployment using Bicep
- Customized Security Hub as a single interface dashboard for reviewing security incidents
- Worked on the IaC inner source Azure shell for automating the infrastructure deployment
- Expertise in Architecting and Implementing Azure Service Offering, such as Azure cloud services, Azure storage, Azure Active Directory (AD), **Azure Resource Manager (ARM)**, Azure Storage, Azure, Blob Storage, Azure VMs, **SQLDatabase**, Azure Functions, Azure Service Fabric, Azure Monitor, and Azure Service Bus.
- Responsible for design and maintenance of the GIT Repositories and the access control strategies, and performed all necessary day-to-day GIT support for different projects
- Set up the CI/CD Pipeline using GitHub action in the SFSBI application
- Automated the Ingress kured and pod-identity deployments into the AKS cluster using Helm and Terraform
- Experienced in Docker containers, writing Docker files, Docker compose, Docker network, and private Docker repository
- Using Jenkins as a **CI/CD** pipeline along with the Developers team.
- Built and deployed Docker containers to break up the monolithic app into micro services, improving developer workflow, increasing scalability, and optimizing speed
- Customized Security Hub as a single interface dashboard for reviewing security incidents.
- Automated the Azure alerts Deployment using Terraform and Azure DevOps pipelines
- Designed, deployed, and managed **REST** and HTTP APIs using AWS API Gateway integrated with AWS Lambda and backend services
- Implemented authentication and authorization for APIs using API Gateway with Cognito User Pools, Lambda

authorizers, and IAM roles

- Created reusable Terraform modules to define **REST API** resources, methods, and integrations with Lambda functions or HTTP backends
- Version-controlled API Gateway configurations using Infrastructure as Code and integrated them into CI/CD pipelines
- Integrated **AWS Security Hub** with Service Now for Incident Management by leveraging Lambda.
- Configured SAML-based authentication by integrating applications in Azure AD
- Configured Azure App Services for hosting .NET applications
- Configured Azure SQL Databases for application management
- Developed Templates for AWS infrastructure as code using **Terraform** to build lower environments, staging, and production environments
- Developed Templates for Azure infrastructure as code using **Terraform** to build environments
- Responsible for managing Jenkins system configuration, plugin management, and RBAC to Jenkins Jobs
- Develop a monthly Cloud Economics report for tracking and forecasting monthly Cloud usage bills and cost optimization
- Developed Templates for AWS infrastructure as code using **Terraform** to build lower environments, staging, and production environments
- Experience in setting up **CI/CD** pipeline integrating various tools with Cloud Bees Jenkins to build and run **Terraform** jobs to create infrastructure in AWS
- Used Jenkins for Continuous Integration and deployment into Tomcat/WebLogic Application Server
- Implemented continuous integration and tracking source code changes using Jenkins