

Sai Krishna

Senior Cloud, Data & AI Platform Engineer

saikrishna7215@gmail.com | +1 937-900-1704 | Denver, CO

PROFESSIONAL SUMMARY

AWS Certified Solutions Architect and Senior Cloud, Data & AI Platform Engineer with 9+ years of experience designing, automating, and scaling cloud-native infrastructure, data platforms, and ML pipelines across **AWS, Azure, and GCP**. Proven expertise in Kubernetes/EKS, Terraform, CI/CD, PostgreSQL, observability, and production MLOps, with measurable results in release velocity, reliability, provisioning speed, and cloud cost optimization.

At Charter Communications, **led large-scale AWS and Kubernetes platform modernization**, building highly available, secure, and cost-optimized environments serving thousands of workloads. Designed end-to-end MLOps workflows with Kubeflow, MLflow, Airflow, and SageMaker, and partnered with GPT/LLM developers to operationalize AI workloads on EKS and SageMaker with CI/CD, observability, auto-scaling, model monitoring, and drift detection.

Recognized for architectural leadership and mentorship across US and India teams, with an **end-to-end engineering perspective spanning cloud operations, data, model delivery, and AI application infrastructure**. Currently expanding hands-on capabilities in Generative AI and agentic architecture, including embeddings, vector search, RAG, LangChain, LlamaIndex, LangGraph, Model Context Protocol (MCP), tool-using agents, multi-agent patterns, and n8n automation.

Experienced across telecom, SaaS, and financial domains, delivering platforms that improved developer productivity, accelerated time-to-market, and reduced cloud spend by 20%+. Combines **automation-first operations, governance-as-code, security, observability, and platform engineering** to move intelligent systems from infrastructure and data foundations through reliable production delivery.

SKILLS

Cloud Platforms: AWS (EC2, Lambda, EKS, RDS/Postgres, S3, SageMaker, Step Functions, EventBridge, Route53), Azure, GCP

Containers & Orchestration: Kubernetes/EKS, Docker, Helm, ArgoCD, OpenShift (exposure)

Infrastructure as Code: Terraform, Ansible, Chef, Puppet, CloudFormation

CI/CD & DevOps Tools: GitLab CI/CD, Jenkins, Bitbucket, GitHub Actions, Packer, SonarQube, GitOps workflows

Programming & Scripting: Python, Bash, Ruby, YAML, JSON

MLOps & Model Operations: Kubeflow, MLflow, Airflow, SageMaker, TensorFlow Serving, TorchServe, model versioning, serving, monitoring, drift detection, automated retraining

Generative & Agentic AI (developing): OpenAI API, Hugging Face, transformers, prompt engineering, embeddings, vector databases (FAISS, Pinecone), RAG, LangChain, LlamaIndex, LangGraph, MCP, tool-using agents, multi-agent patterns, n8n

Databases: PostgreSQL (RDS/Postgres), MySQL, MongoDB, Cassandra, DynamoDB

Monitoring & Observability: DataDog, Prometheus, Grafana, ELK (Elasticsearch, Logstash, Kibana), Splunk, CloudWatch

Testing & QA: JMeter, Postman, RestAssured, Selenium (exposure)

Security & Compliance: IAM, InSpec, SAST, Secrets Management, Encryption (KMS, TLS)

Collaboration & Project Management: JIRA, Agile/Scrum, Confluence

GENERATIVE AI & AGENTIC AI DEVELOPMENT

- Structured professional development in Python-based Generative AI, NLP and deep learning foundations, transformers, LLM application design, embeddings, vector databases, RAG, fine-tuning concepts, evaluation, and deployment patterns.
- Current focus: LangChain/LlamaIndex orchestration, LangGraph stateful agents, MCP-based tool integration, memory and planning, multi-agent design, and n8n workflow automation—grounded in production practices for security, CI/CD, observability, and cloud scale.

EXPERIENCE

Senior DevOps/MLOps Engineer, Denver, CO

Charter Communications, June 2022-Present

- Designed & deployed **enterprise-grade AWS infrastructure** (EC2, S3, RDS, Lambda, VPC, ELB, Route53) with HA/DR strategies, supporting thousands of internal and customer-facing workloads.
- Built and **optimized multi-stage CI/CD pipelines** (Jenkins, GitLab, Bitbucket) that reduced release cycle time from 3 days to under 6 hours.
- Migrated core workloads to **Kubernetes/EKS** with Helm & ArgoCD GitOps workflows, improving release reliability by 40%.
- Implemented **Terraform modules** and standardized infra automation across environments, **cutting provisioning times by 70%**.
- Developed and enforced **security/compliance guardrails** (InSpec, IAM policies, encryption, tagging standards) for systems handling PII and regulated data.
- Designed and automated **end-to-end MLOps pipelines** (data preprocessing, model training, evaluation, and deployment) using tools like Kubeflow, MLflow, and Airflow.
- **Implemented MLOps practices by extending CI/CD pipelines to ML workflows**, automating data

preprocessing, model training, validation, and deployment to **AWS EKS & SageMaker**.

- Implemented CI/CD pipelines for **ML models integrating GitLab/GitHub Actions** with containerized environments
(Docker + Kubernetes).
- Automated **model retraining workflows** triggered by new data/events using Step Functions, EventBridge, and Airflow.
- Partnered with GPT/LLM developers to operationalize **internally built GPT applications** on AWS EKS/SageMaker, integrating CI/CD pipelines, monitoring, and scaling for inference workloads.
- **Built scalable model serving infrastructure** with AWS SageMaker / TensorFlow Serving / TorchServe, ensuring low-latency inference.
- Established model monitoring and drift detection (data drift, concept drift, performance degradation) with
Prometheus, Grafana, and custom alerts.
- Automated AMI and container image pipelines with **Jenkins + HashiCorp Packer**, streamlining immutable infrastructure adoption.
- Automated model retraining workflows triggered by new data/events using **AWS Step Functions / EventBridge / Airflow**.
- Integrated **SonarQube & SAST tools into CI/CD pipelines**, raising code quality **scores by 30%**.
- Championed cloud cost optimization, leading initiatives that **saved the org 20%+ in monthly AWS spend**.
- **Mentored 5 junior engineers** and collaborated across US and India teams to standardize **GitOps practices and Terraform workflows**.
- Experienced working with **Azure and GCP** using the terraform by automating the resource creation seamless across the environments.
- Profusely worked on Version control tool GIT to track changes made by different people in source code.
Performed Branching, Tagging, Release Activities on Version Control Tools like **GIT and Bit Bucket & GitLab**.
- Working on **Ruby scripting** to automate the framework that reaches out to AWS services via the AWS CLI libraries to run scripts and commands.
- Developing scripts for deployment of customer environments into **AWS using Bash, Python** and created scripts in Python which integrated with Amazon API to control instance operations.
- **Responsible for Continuous Integration and Continuous Delivery (CI/CD)** process implementation using

Jenkins along with Python and Shell scripts to automate routine jobs.
- Automate and deploy development environments using Ansible to various cloud providers **Amazon AWS and EC2, Rackspace, Azure**. Travis and Jenkins testing.
- Hands-on with **PostgreSQL administration (RDS/Postgres)** including schema tuning, query

optimization, and backup automation.

- **Configuring Sharding, monitoring**, and identify the proper shard key Selection, inserting large amounts of data Performing Security measures, backups and **restore backups for MongoDB**.
- Written **Chef Cookbooks** and recipes in Ruby to Provision several pre-prod environments consisting of **Cassandra DB installations**, WebLogic domain creations and several proprietary middleware installations.
- **Supported Deployments into PROD & Pre-Prod environments** with multiple Application server technologies like WebLogic, JBoss and Apache Tomcat.
- **Extensive experience in JIRA** for creating bugs tickets, storyboarding, pulling reports from dashboard, creating and planning sprints.
- Extensively worked on **JMeter** to create Thread Groups and test Web Application for various loads on key business scenarios.

Key Tech: AWS (EKS, Lambda, EC2, RDS/Postgres, SageMaker, VPC, Route53), Kubernetes, Helm, ArgoCD, Terraform, GitLab CI/CD, Jenkins, Python, Bash, Ruby, MLflow, Kubeflow, Airflow, ELK, Splunk, DataDog, Prometheus, Grafana, Ansible, Chef, Puppet.

DEVOPS ENGINEER, HYDERABAD, India

Kantar, January 2020-July 2021

- Automated infrastructure provisioning with **Terraform & Ansible**, improving release **speed by 60%**.
- Deployed centralized logging with **ELK + Fluentd**, enabling proactive anomaly detection across distributed apps.
- Administered CI/CD pipelines with **Jenkins & SonarQube**, enforcing code quality and security standards.
- Conducted **performance testing (JMeter)** and integrated it into release cycles for high-scale applications.
- Maintained source control with **Git & Bitbucket**, enabling efficient branching and release management.
- Implemented a continuous delivery pipeline with **Docker, Jenkins and GitHub** and AWS AMI's whenever a new GitHub branch gets started. Jenkins, our continuous integration server, automatically attempts to build a new Docker continuous from it.
- Automated various infrastructure activities like Continuous Deployment, Application Server setup, Stack monitoring using Ansible playbooks and had Integrated **Ansible with Rundeck and Jenkins**.
- Involved in Jira as defect tracking system and configure various workflows, customizations, Plug-ins for Jira bug/issue tracker and, also involved in administration and support of **Unix Servers including SUSE, Solaris, Red Hat Linux and CentOS**.
- Responsible for monitoring servers using **Cloud watch and using ELK Stack, Elastic Search, Logstash and Kibana**, developed a RESTful API of Elasticsearch to analyze, search and visualize data.
- Setup and maintained automated environment using **Chef Recipes within Azure environment**. Good understanding of Knife, **Chef Bootstrap process etc**.
- Configured Centralized log collection using **Fluentd, Logstash, Elasticsearch and Kibana (ELK) stack** for parsing common web server, app server logs, Syslog for intrusion detection.

- Maintain **GIT, Bitbucket repository**, handling branching, Merging, tagging and release activities.
- Used **AWS Beanstalk** for deploying and scaling web applications and services developed with **Java, PHP, Node.js, Python, Ruby, and Docker** on familiar servers such as Apache, and IIS.

DEVOPS ENGINEER, Hyderabad, India

MetrixLab, August 2017-December 2019

- Migrated **50+ legacy servers from on-prem to AWS Cloud**, improving resilience and scalability.
- Managed infrastructure automation with **Puppet & Chef**, ensuring consistent provisioning across multiple environments.
- Built Jenkins pipelines for automated builds/deployments, cutting release times in half.
- Implemented Docker-based containerization, reducing environment **setup time by 70%**.
- Monitored infrastructure with **Nagios & CloudWatch**, reducing **downtime by 20%**.
- **Managed build & release cycles with SVN, Ant, Jenkins**, supporting multiple parallel projects.
- Supported multi-stage environments (**DEV, QA, UAT**), ensuring consistent deployments.
- User Management - creating users, assigning roles, managing permissions for **SQL Server, Mongo DB database server**.
- Configured Ansible to manage AWS environments and automate the build process for core AMIs used by all application deployments including Auto scaling, and Cloud formation scripts.
- Implemented a continuous delivery pipeline with **Docker, Jenkins and GitHub and AWS AMI's** whenever a new GitHub branch gets started. Jenkins, our continuous integration server, automatically attempts to build a new Docker continuous from it.
- Participated in the release cycle of the product which involves environment like **development QA and UAT**.
- **Used Apache ANT to prepare and build war files for java-based web applications**. Developed ANT scripts to prepare build.xml files to the applications.
- Resolved merging issues during build and release by conducting meetings with developers and managers.
- **Involved in editing the existing ANT (build.xml) files** in case of errors or changes in the project requirements.

EDUCATION

MASTER OF SCIENCE (M.S.) IN COMPUTER SCIENCE, Dayton, OH, US

Wright State University, May 2023

AWARDS

Rising Star of the Year (2019); Best Employee of the Year (2021); Best Team Player of the Year (2023)