

SHUBHAM PAI

(860) 906-2294 | shubham.pai@uconn.edu | [linkedin.com/in/shubhampai](https://www.linkedin.com/in/shubhampai) | [shubhampai99.github.io](https://github.com/shubhampai99)

SKILLS

Programming: Terraform, Python, Bash, JSON, YAML

Containerization: Rancher Kubernetes Engine (on-prem), EKS, ECS, Docker, Helm, kubectl, k9s

SDLC/Automation: Git, GitLab CICD, GitHub Actions, BitBucket, Jenkins, ArgoCD, SonarQube, Snyk, Nexus, ECR

Platforms: AWS, VSphere, Cloudflare, RHEL

Service Management: ServiceNow, Jira, Confluence

Certifications: AWS Certified Cloud Practitioner (May 2022)

EXPERIENCE

IdenTrust | DevOps Engineer

June 2024 – Present

- Operating in small, fast-paced teams as a lead DevOps engineer for the IdenTrust and HydrantID organizations, spearheading best practices and implementing developer requests within Kubernetes and AWS ECS environments.
- Implemented a self-service process to enable developers to deploy versioned helm charts to various kubernetes environments using Gitlab CI, JSON, Python, Terraform, and Rancher Continuous Delivery. This process standardized repository settings, branching patterns, and versioned charts for production deployments.
- Conducted various upgrades on Rancher, Kubernetes and EKS, ensuring critical uptime and providing services to developers and customers.
- Structured repositories and produced automation for AWS services using Terraform, GitLab CI, and Bash.
- Enabled and integrated alerting for applications in AWS and Kubernetes to MS Teams and PagerDuty to provide better time-to-response for application failures.
- Sponsored efforts and produced prototypes to expand IaC use-cases among Ops teams to establish stateful, version-controlled infrastructure definitions.
- Created a dashboard from Snyk scans for vulnerability scanning of Kubernetes management-plane applications
- Guided workshops to teach the operations organization on containerization, kubernetes, CICD, and other DevOps principles.

The Cigna Group/Evernorth Health Services | TECDP Lead Analyst

June 2021 – May 2024

Digital DevSecOps Engineer:

- Built, managed, deployed, and troubleshot multi-region AWS cloud infrastructure using Terraform, ensuring compliance with enterprise governance and allowing delivery partners to maintain focus on their key competencies.
- Created Elastic Kubernetes Service (EKS) deployments and self-service pipelines to accelerate the delivery of Angular, NodeJS, Java, and Go applications using Helm, Terraform, and GitLab CICD.
- Implemented the Cloudability metrics agent controller into EKS clusters using Helm and Terraform to gain visibility on cluster usage and guide cost-saving directives.
- Outfitted a notification step into digital deployment pipelines to notify production support of changes being actively deployed into production environments using Gitlab CICD, YAML, and Bash.
- Instrumented a process to replicate AWS S3 buckets to other AWS accounts using Terraform while providing an easy to-use JSON interface, allowing developers to easily enable replication and configure destinations.

Data Quality and Tools Software Engineer:

- Built a deployment pipeline to deploy HQL files to Hadoop, automate testing, and safeguard production workflows from failures using Conduit (in-house pipeline language), Bash, Python, and UDeploy.
- Utilized object-oriented programming and multi-threading in Python to generate CSV files as part of a departmental migration from an Oracle database to a REST API as a source of reference data.
- Modified database security schemas using SQL and Python to ensure cloud data could be viewed in Teradata to unify data access and enable security views across the organization.
- Migrated code deployment pipelines using Bash and Python as part of an organization-wide migration from Jenkins to CloudBees, resulting in a smooth and successful transition for the developer team.

Big Data Ingestion Software Engineer:

- Architected and executed a migration of code bases from GitLab to GitHub, which involved refactoring Python APIs, managing repositories' branch rules and webhooks, and updating security tokens in infrastructure; Concluded the migration, ensuring critical uptime, seamless transition, and 0 application blackouts for end users.
- Integrated unit tests and environmental checks to pass CI/CD pipeline standards so code can be certified for deployment into lower and upper environments.
- Created a process to move data from tables in Hadoop to an AWS S3 bucket using Python, Bash, AWS Athena, and AWS Glue as part of a temporary data storage solution in the cloud.
- Enabled a workflow for data engineers to utilize ingestion of a subset of tables by allowing a custom WHERE clause on incremental data loads, reducing workflow execution times and retaining data quality for downstream users.

The Cigna Group/Evernorth Health Services | Big Data Infrastructure Engineer Intern May 2020 – May 2021

- Utilized the Cloudera Manager (CM) python API to get updated schemas and roles on servers, services, and environments in CM
- Produced automation to generate ansible host files for upper and lower environments to reduce overhead and give team members the most up-to-date roles for an environment for their automation.
- Created a dashboard for organizing information on 1,000+ Hadoop servers to coordinate and assist with RHEL upgrades on environments using CM API, HDFS, Hive, and Looker

EDUCATION

University of Connecticut | Storrs, CT

Bachelor of Science in Engineering, May 2021

Major in Computer Science and Engineering; Minor in Ornamental Horticulture

Honors: Cum Laude