

# SAMUEL RAJU MAMOOTIL

## DevOps Engineer

Austin, TX | 512.563.0303 | samuelrajumamootil@gmail.com | linkedin.com/in/samuel-mamootil | samuelmamootil.github.io

### PROFESSIONAL SUMMARY

---

DevOps Engineer with 7+ years of professional experience helping enterprise teams modernize applications, automate cloud delivery, and operate secure platforms across AWS, Azure, and GCP. Most recently worked in an AWS data engineering platform for client Alight Solutions, building Terraform-driven infrastructure and Python automation for data ingestion, transformation, metadata discovery, analytics delivery, and forecasting workflows. Previously with IBM delivered migration and DevSecOps work for clients including H&M, Siemens, Südzucker, Parexel, GenRe, and Veolia Water, including 100+ application migrations, CI/CD pipeline design, hardened images, private connectivity, monitoring, and technical documentation.

### TECHNICAL SKILLS

---

**Cloud Platforms:** AWS, Azure, GCP

**Infrastructure as Code:** Terraform, ARM Templates, Bicep, CloudFormation

**Scripting and Development:** Python, Bash, PowerShell, PHP, SQL, PySpark

**Monitoring and Documentation:** Azure Monitor, Application Insights, CloudWatch, CloudTrail, Prometheus, DataDog, Confluence

**Data Engineering and Analytics:** AWS Glue, Athena, Redshift, Lake Formation, SageMaker, Azure Data Factory, Databricks, Azure Synapse, Hive, PostgreSQL, Tableau, Power BI

**Security and Quality:** SonarQube, Coverity, IAM, KMS, Secrets Manager, Azure Key Vault, CloudTrail

**CI/CD and Containers:** GitHub Actions, Azure DevOps YAML, GitLab CI/CD, Docker, Kubernetes, AKS, Helm charts

### PROFESSIONAL EXPERIENCE

---

#### Wipro | DevOps Engineer - AWS Data Platform

Client: **Alight Solutions (USA/AWS)** | Navi Mumbai, India | February 2024 - April 2025

- Led and mentored a 7-member data engineering team; wrote, deployed, and managed reusable Terraform modules to provision AWS data platform components for ingestion, transformation, cataloging, analytics, caching, and ML-ready workflows across S3, Glue, Athena, Redshift, Lambda, Step Functions, and Redis.
- Worked in a data engineering workflow where data from 10+ applications landed in Amazon S3; configured AWS Glue Crawlers to scan new and changed datasets, infer schemas, update partitions, and register table metadata in AWS Glue Data Catalog; and used Athena and Redshift SQL to validate, join, and analyze curated datasets consumed by Tableau dashboards.
- Built Glue/PySpark and SQL-based validation checks to reconcile source-to-target record counts, detect nulls and duplicate business keys, validate schemas, date formats, accepted value ranges, and identify sensitive fields before publishing curated datasets.
- Applied access controls and data protection practices using Lake Formation, IAM, KMS, and Secrets Manager to secure data pipelines, encrypt data, manage credentials, and restrict sensitive datasets at the database, table, row, and column level.
- Wrote custom PII scanning process using Python, AWS Lambda, and EventBridge Scheduler to regularly inspect S3 data for sensitive information such as names, email addresses, phone numbers, credit card numbers, Social Security numbers, and exposed credentials.
- Deployed and configured CloudWatch and CloudTrail to monitor pipeline execution, review job failures, track access activity, and support auditability across Glue, S3, Athena, Redshift, and security-related services.
- Supported Redshift as both a producer and consumer data platform, providing governed access to internal users and enabling other teams to consume trusted datasets for reporting and downstream processing.
- Developed Python batch jobs scheduled to run daily at midnight to fetch the latest source data, process updates, and publish refreshed datasets into Redshift for downstream reporting and Tableau dashboard refreshes.
- Prepared curated datasets in S3 and Redshift for consumption by other application teams, advanced analytics workflows, and machine learning use cases, including Amazon SageMaker model training and prediction.
- Worked with application developers and product owners, to gather requirements, understand data definitions, document ownership, clarify access needs, and deliver datasets for reporting, governance, and analytics.

#### IBM | Junior Tech Lead - DevOps

Clients: **H&M (Sweden/Azure), Siemens (Germany/AWS), Nokia, Südzucker (Germany/GCP)** | Navi Mumbai, India | November 2021 - February 2024

- Led multi-client DevOps delivery across Azure, AWS, and GCP, wrote Terraform, GitHub Actions, Azure DevOps, Helm, Python, Bash, hardened images, migration planning, and secure release automation.
- For H&M, planned and architected GitHub Actions and reusable Terraform modules, migrating 100+ legacy applications and automating VM, AKS, Azure Web App, and Azure Function deployments.
- Deployed and maintained Azure Storage Accounts for Terraform remote state, enabled versioning, and implemented private endpoints to secure internal application communication.
- Coordinated with developers, product owners, support teams, and stakeholders to migrate applications based on target

architecture (Lift and Shift, Re-architecture, Modernization) and remediate deployments across AKS, Web Apps, and Functions with scaling enabled.

- Strengthened Azure infrastructure and application security using encryption at rest/in transit, Key Vault based secret retrieval during deployment, Azure AD application setup, and hardened deployment patterns.
- Automated dynamic allowlisting and removal of GitHub Actions public IP addresses during deployments and wrote Python/Bash scripts to support secure, repeatable release execution.
- Implemented DevSecOps quality gates with SonarQube, Coverity, and secrets hardcoding /vulnerability checks, failing builds on known high vulnerabilities and notifying application teams through Microsoft Teams channels.
- For Siemens, modernized highly dependent legacy applications using custom base images, Helm charts, terraform modules, and repeatable deployment using GitHub action for application build and AWS resources deployment using OIDC connection.
- Documented architecture, delivery decisions, SOPs, SLOs, SKA/knowledge artifacts, and handover materials in Confluence to showcase completed work and support operational readiness.
- For Südzucker, created Azure DevOps OIDC service connections and a GCP deployment proof of concept with Terraform modules for Google resources including GFE/GHE components and storage, then handed over documentation and operating steps.

### **LTIMindtree / LTI | Software Engineer - Cloud and DevOps**

Clients: **Parexel (USA/Azure), GenRe (USA/Azure), Veolia Water (AWS)** | Bengaluru, India | April 2018 - November 2021

- Delivered cloud automation for Parexel, GenRe, and Veolia Water, supporting landing zone provisioning, application deployments, networking, identity, storage, security, and release operations across Azure, AWS, and GCP.
- For Parexel, wrote ARM Templates, Terraform, Bicep, Azure DevOps YAML, PowerShell, Bash, and Python automation to deploy Azure compute, networking, identity, and storage services.
- Patched Windows and Linux Azure VMs, created hardened VM images, and published standardized images through Azure Compute Gallery.
- Built Python automation to regularly lock resources and reduce accidental deletion risk in cloud environments.
- Built one-click PHP application for LTI Cloud Practice, provisioning landing zone infrastructure across Azure, AWS, and GCP from standardized user inputs and SQL for database and terraform for deployment.
- Deployed and managed Azure data platform components including Azure Blob Storage, ADLS Gen2, Azure Data Factory, ADF pipelines, Azure Databricks, and PostgreSQL, while supporting deployment/configuration of Hive and Azure Synapse components for data processing, metadata, and analytics workflows.
- Built and scheduled ADF pipelines to ingest data from source systems, copy data from Synapse/PostgreSQL into ADLS Gen2, orchestrate workflow dependencies, trigger Azure Databricks notebooks, and use PySpark, SQL, and Hive-based logic to clean, transform, and publish curated datasets consumed by Power BI dashboards.
- Supported GenRe and Veolia Water engagements with cloud deployment automation, release support, security remediation, and operational readiness patterns.

### **NeuroSpeech Technologies | Junior Software Developer**

Navi Mumbai, India | August 2017 - April 2018

- Wrote Android application features across requirements, implementation, debugging, backend integration, and release support in a startup environment.

### **ADDITIONAL EXPERIENCE**

---

#### **UTA College of Business | Student Web Developer - Part-Time**

Arlington, TX | June 2025 - May 2026

- Maintained production web updates across 4 university platforms, resolving access, content, release, and site issues while coordinating status with business users.

### **KEY PROJECTS**

---

#### **ReRhythm - Generative AI Platform | AWS Bedrock, Claude, ChromaDB, RAG, CloudFormation, ECS**

- Architected, build and wrote a scalable Generative AI platform demonstrating enterprise AI capabilities using AWS Bedrock, Claude, and ChromaDB.
- Deployed a Retrieval-Augmented Generation (RAG) pipeline to ingest, process, and analyze complex document data, enabling the system to identify skill gaps and generate intelligent workflow recommendations.
- Deployed the underlying infrastructure entirely via CloudFormation using ECS and AWS-native platform services.

### **EDUCATION**

---

The University of Texas at Arlington | Master of Science in Information Systems | GPA: 3.7/4.0 | May 2026

The University of Texas at Arlington | Graduate Business Analytics Certificate | GPA: 4.0/4.0 | May 2026

Mumbai University | Bachelor of Engineering in Computer Engineering | May 2017

### **HONORS**

---

- Maverick Advantage Certificate of Distinction, UTA, 2026.
- Second Runner Up, ReRhythm AI Application, AWS Hackathon / UTA Symposium, 2026.
- Best Team Award, Parexel Project, LTIMindtree, 2019.