

# SAGAR BUDHATHOKI

## Senior DevOps / SRE Engineer

Kathmandu, Nepal (remote) · [sagar@sagarbudhathoki.com](mailto:sagar@sagarbudhathoki.com) · [sagarbudhathoki.com](http://sagarbudhathoki.com) · [github.com/sbmagar13](https://github.com/sbmagar13) · [linkedin.com/in/sbmagar13](https://linkedin.com/in/sbmagar13)

### SUMMARY

Senior DevOps / SRE Engineer with 5+ years of experience. Sole platform owner of a Swedish multi-tenant event-management SaaS on AWS: ECS Fargate, Aurora PostgreSQL with schema-per-tenant, ElastiCache Redis, Amazon MQ in eu-north-1. Building AI agents for ops (MCP, LangGraph, local LLM inference); open to remote senior roles.

### EXPERIENCE

**DevOps / SRE Engineer** · Threadcode Technologies Pvt. Ltd. · EventLogic, Swedish multi-tenant event-management SaaS · Lalitpur, Nepal  
May 2023 · present

- Owner of the entire AWS platform end to end: ECS Fargate, Aurora PostgreSQL, ElastiCache Redis, Amazon MQ in eu-north-1.
- Diagnosed and resolved a 19-minute full-platform outage caused by blocking Redis KEYS calls exhausting the Tomcat/JDBC thread pool. Added connection-pool checkout timeouts, tuned RDS parameters, then drove a 68-task reliability program across 11 epics and 7 sprints to prevent recurrence.
- Established a cross-region DR path where none existed: Aurora Global Database from eu-north-1 to eu-west-1, EFS and ECR replication, shared KMS keys, and a documented runbook for promotion.
- One Python API call sets up schema-per-tenant on Aurora, wires SQS and EventBridge, creates ALB listener rules, provisions a CloudFront / S3 distribution, configures Route 53 records, and registers the tenant in DynamoDB.
- Stood up OneUptime on a K3s cluster in a separate region (eu-central-1) for status pages, uptime monitoring, on-call scheduling and incident management. Deliberately off the primary region so observability survives a primary-region outage.
- Self-managed three-node Elasticsearch cluster orchestrated with Terraform and Ansible. Split deploy and split-restart playbooks so a single config change cannot cascade across the cluster.
- Cut monthly AWS spend by removing orphaned NAT Gateways, adding S3 and DynamoDB gateway endpoints to drop data-transfer cost, setting log-retention policies on CloudWatch, and right-sizing EBS volumes.
- Technical Reviewer for Python for DevOps (Packt).

**DevOps Engineer** · Cloudyfox Technology Pvt. Ltd. · Kathmandu, Nepal Sep 2021 · Apr 2023

- Ran Kubernetes for containerized workloads at Cloudyfox. Built CI/CD on GitLab CI and Jenkins for both app and infra deploys.
- First Terraform / Terragrunt at scale across AWS.
- SysOps, Linux admin, OpenVPN, centralized logging with CloudWatch, ELK, Grafana.

**AI/ML and Backend Engineer (earlier roles)** · VolgAI · Genese Cloud Academy · IBZ Networks · Kathmandu, Nepal 2020 · 2021

- Built AI chatbots with RASA (NLP), backend APIs with Django and Flask, RTSP/FFmpeg pipelines for CCTV image processing.
- Async work via Celery and RabbitMQ. AWS AI/ML Internship at Genese.

### SKILLS

**Cloud & Infrastructure:** AWS, ECS Fargate, AWS Lambda, Amazon MQ, CloudFront, S3, API Gateway, Route 53, VPC Networking, EFS, ECR, Terraform, Terragrunt, CloudFormation, AWS CDK, Ansible, Docker, **Kubernetes (K3s)**

**CI/CD:** GitLab CI, Jenkins, CodePipeline, GitHub Actions

**Observability:** Prometheus, Grafana, Loki, OpenTelemetry, AWS CloudWatch, OneUptime, ELK Stack, Elasticsearch

**Databases:** Aurora PostgreSQL, ElastiCache (Redis), DynamoDB, AWS DMS, PostgreSQL, Redis

**Languages & Frameworks:** Python, FastAPI, Django, Flask, Bash, JavaScript

**Security:** AWS IAM, AWS KMS, Secrets Manager, AWS Inspector, CloudTrail, OpenVPN

**AI / ML:** Anthropic MCP, LangGraph, LangChain, Local LLM Inference, PyTorch, RASA

**OS & Tooling:** Apache Airflow, Airbyte, Celery + RabbitMQ, FFmpeg, Arch Linux, Ubuntu

### PROJECTS

**hashnode-mcp-server** · AI assistant integration · [github.com/sbmagar13/hashnode-mcp-server](https://github.com/sbmagar13/hashnode-mcp-server)

Open-source Model Context Protocol server that wires AI assistants like Claude into the Hashnode content API. The pattern carries into the broader agentic-DevOps work.

**VQGAN-CLIP-Text-to-Image** · text-to-image · 2021 · [github.com/sbmagar13/VQGAN-CLIP-Text-to-Image](https://github.com/sbmagar13/VQGAN-CLIP-Text-to-Image)

Multimodal text-to-image generation using VQGAN + CLIP architectures in PyTorch.

### CERTIFICATIONS & RECOGNITION

Technical Reviewer, Python for DevOps · Packt Publishing 2023

Generative AI: From GANs to CLIP with Python and PyTorch · Udemy 2021

AWS AI/ML Internship · Genese Cloud Academy 2020 · 2021

### EDUCATION

Bachelor in Computer Engineering · Tribhuvan University, Western Regional Campus (IOE) · Pokhara, Nepal 2016 · 2020