

Satyam Kumar

✉ ksatyam0468@gmail.com ☎ 6296178009 📍 Bengaluru, Karnataka 🌐 LinkedIn 📁 Portfolio

Skills

- **Containerization & Orchestration:** Kubernetes (**Kind**, Helm Charts), Docker, Docker Compose.
- **CI/CD & Observability:** Git, **GitHub Actions**, **Jenkins** (Pipelines), **Prometheus**, **Grafana**, n8n Workflow Automation.
- **Cloud & IaC:** Google Cloud Platform (GCE, GCS, Cloud Functions), AWS (EC2), **Terraform**, Cloudflare Zero Trust.
- **Operating Systems & Storage:** Linux (Ubuntu), TrueNAS Core (ZFS).
- **Programming & Automation:** Python, SQL, YAML, Bash.
- **AI & Operations:** RAG Pipelines, Vector Databases (Qdrant, Milvus).
- **AI-Native Development:** Model Context Protocol (MCP), Agentic Workflows (Gemini CLI, Claude Code).

Projects

Kubernetes CI/CD & Observability Pipeline

Feb 2026

- **Cluster Orchestration:** Provisioned a multi-node Kubernetes cluster using **Kind** to simulate a production-grade container environment on local Linux infrastructure.
- **CI/CD Automation:** Configured a **Jenkins** pipeline to automate the software delivery lifecycle, setting up triggers to build and push Docker images upon code commitment.
- **Infrastructure Monitoring:** Deployed **Prometheus** to scrape real-time cluster metrics and node health data, ensuring high availability of deployed services.
- **Data Visualization:** Designed custom **Grafana** dashboards to visualize CPU usage, memory consumption, and network traffic.

Cloud-Native RAG Application (GCP)

Jul 2025 – Sep 2025

- **Serverless Architecture:** Architected an event-driven pipeline on Google Cloud Platform to process unstructured data, utilizing Cloud Functions for automated triggers and GCS for object storage.
- **Microservices Deployment:** Engineered a containerized stack using Docker Compose to orchestrate Milvus (Vector DB), Ollama (Model Serving), and a FastAPI backend.
- **Security & Networking:** Deployed infrastructure on Google Compute Engine (GCE), configuring strict firewall rules for secure ingress and optimizing VM resources for quantized model performance.
- **Pipeline Automation:** Designed automated ingestion scripts to parse, chunk, and embed data automatically upon upload.

NLP Inference Service Deployment (FactChecker) 📁

Oct 2024 – Mar 2025

- **Cloud Migration & Compute:** Migrated training workflows to Google Colab (T4 GPU) to accelerate processing cycles and resolve local hardware limitations.
- **Containerized Deployment:** Dockerized a Python-based asynchronous inference API (Flask) and deployed it to Hugging Face Spaces, ensuring a consistent runtime environment.
- **Resource Optimization:** Optimized backend service dependencies to decrease memory footprint by 40% while maintaining 97% functional API accuracy.

Private Cloud & Network Security Administration

May 2020 – Oct 2025

- **Zero Trust Security:** Implemented Cloudflare Zero Trust tunnels to securely expose local services to the public internet without opening router ports, defining strict access control policies.
- **Storage Management:** Constructed and managed a Network Attached Storage (NAS) solution using TrueNAS Core, configuring ZFS storage pools for data integrity and redundancy.

Education

BSc(Hons) Computer Science, Goldsmiths, University of London

Apr 2021 – Mar 2025 | London, UK

Certificates

Google IT Support Professional 📁 — Professional certificate by **Google** via Coursera

Intro to Quantum Computing 📁 — Qubit x Qubit (in partnership with **IBM Quantum**)