

Shivam Chaubey

Backend & Infrastructure Engineer

Mumbai, India | +91-9321664428 | shivamchaubey027@gmail.com | [GitHub](#) | [LinkedIn](#) | [Portfolio](#)

PROFESSIONAL SUMMARY

Backend Engineer building search and discovery infrastructure for Lykstage at Manhattan Tech Ventures. Specialized in Go microservices, Kubernetes orchestration, and AWS cloud architecture. Track record of reducing system RTOs from hours to <15 minutes and optimizing data pipelines for 30K+ entity processing

SKILLS

Backend Engineering: Go, Python, REST APIs, Microservices, Goroutines, PostgreSQL, Clickhouse, MySQL

Cloud & Orchestration: AWS (VPC, EC2, RDS), Kubernetes (Client-go, Operators, RBAC), Terraform, Docker

Observability Stack: Prometheus, Grafana, Custom Metrics Exporters, PromQL, Time-Series Databases

DevOps & Tooling: GitHub Actions (CI/CD), Linux Internals, Git, Multi-Stage Docker Builds

EXPERIENCE

Manhattan Tech Ventures | Backend Engineer(Go) – Remote Jan 2026 – Present

- Architected a Vespa-based discovery pipeline indexing 34K+ creator and video entities using a custom doctype discriminator; reduced processing time by 87% via Redis deduplication (4K incremental vs 30K full scans)
- Engineered an **async engagement** buffer using **RabbitMQ** and feature-flagged Go handlers, cutting like/dislike write latency from **4s** → **50ms** by **decoupling synchronous MySQL writes** into domain-driven event queues (queue.events.like, queue.events.dislike)
- Designed a domain-driven event handler registry across 3 packages (engagement, core, safety) with config-driven queue initialization via YAML, replacing a monolithic definitions.go to support RFC-aligned **scaling to 15+ queues**
- Built a Go transformation service mapping legacy MySQL user records to Vespa "channel" entities, resolving SQLC generation issues and field-mapping conflicts to achieve **100% data consistency** between video and channel search results

PROJECTS

Kubernetes Monitoring & Autoscaling System | Go, Prometheus, CRDs, K8s [Link](#)

- Architected a custom **Kubernetes autoscaling controller in Go** that scales deployments based on real-time application metrics (RPS, queue depth) rather than static CPU/Memory limits.
- Engineered **'WatcherBot', a custom metrics exporter** to track active task queues and business-critical **SLIs** every 15 seconds.
- Designed complete **Terraform IaC manifests** managing the full monitoring stack (Prometheus, Grafana, Node Exporter, Kube-State-Metrics) with proper RBAC and persistent storage configuration.

Multi-Region Disaster Recovery Platform | AWS, Terraform, RDS [Link](#)

- Designed a **resilient multi-region architecture** on AWS using **Terraform**, reducing RTO (Recovery Time Objective) from hours to **<15 minutes**.
- Implemented **database replication strategies** (RDS Read Replicas) to guarantee **99.99% data durability** across geographically isolated availability zones (us-east-1, eu-north-1).
- Automated infrastructure state management and failover procedures, documenting cost-availability trade-offs in technical decision logs.

Kubernetes Resource Monitor & Cost Estimator | Go, K8s Client-go, PostgreSQL [Link](#)

- Engineered a high-throughput **K8s observability microservice in Go**, utilizing **Goroutines** to process concurrent pod metrics and ensure real-time cluster visibility.
- Exposed actionable insights via a **RESTful API (Gin)**, optimized for low-latency queries by implementing **PostgreSQL** time-series data storage.
- Integrated **AWS Pricing Models** dynamically to calculate cloud spend variance based on fluctuating CPU/Memory requests.

EDUCATION

University of Mumbai Sept 2022 – June 2026

B.E. in CSE (AI & ML) Mumbai, India

- Relevant Coursework:** Distributed Systems, Operating Systems, Data Structures & Algorithms, Computer Networks, Database Management, Cloud Computing