

MUKKTINAADH RAGHAVARAPU

+91-9676349758 | mukktinaadh@gmail.com | linkedin.com/in/mukktinaadh | github.com/mukktinaadh | rmnportfolio.vercel.app

EDUCATION

Parul University

B.Tech Computer Science Engineering — Big Data Analytics

Vadodara, Gujarat

2022 – 2026

- Relevant Coursework: Distributed Systems, Database Systems, Cloud Computing, Data Structures & Algorithms

PROFESSIONAL EXPERIENCE

Software Engineer Intern — AI Product

Feb 2026 – May 2026

MyOnsiteHealthcare

Onsite

- Designed and shipped a production AI scheduling system end-to-end — RAG pipeline (LangChain + vector search) over patient and staff data auto-assigns field workers to visits; live in production at **90%+ accuracy**, cutting manual dispatcher workload by **~60%** across 300+ daily visit assignments on FastAPI + PostgreSQL + Redis + AWS.
- Built multimodal-compatible AI agent executing natural language commands (override assignments, reschedule visits, update patient data, answer queries) — integrated end-to-end with full backend stack.
- Engineered agentic CI/CD system: Cursor AI ingests user script + pipeline template, adapts pipeline via GitLab MCP, pushes code, triggers runs, auto-rolls back failures, and retries with fixes — zero manual intervention.

Volunteer Software Developer

Oct 2024 – Sep 2025

Google Developer Group (GDG)

Remote

- Built Java/Python backend tooling adopted by **200+ students** across 5+ hackathons; mentored 15+ developers on system design, backend architecture, and CI/CD practices.

FEATURED PROJECTS

PyroQueue | Python, Redis, FastAPI, Prometheus, Grafana, Docker

Jun 2026

- Built distributed task queue from scratch — priority queues via Redis sorted sets ($O(\log N)$ insert, $O(1)$ pop-min), worker pool via Python multiprocessing, exponential backoff retry (3x), and dead-letter queue; **no Celery or RQ**.
- Load tested at **893 req/sec** with **p99 <170ms** under 100 concurrent workers, 0% failure rate across 60s sustained load; instrumented with Prometheus + Grafana (queue depth, throughput, failure rate).

FinAI Advisor Enterprise | Python, FastAPI, PostgreSQL, Redis, Kafka, WebSocket, AWS, Docker

Nov 2025

- Architected event-driven microservices platform supporting **1,000+ concurrent WebSocket connections** (load tested with Locust); p99 latency <100ms, PostgreSQL read-replicas + Redis cache at 78% hit rate cut DB load by 60%.
- Implemented JWT/OAuth2 auth + API gateway rate limiting (token bucket, 1,000+ req/min); deployed Prometheus + Grafana observability stack with automated alerting across all microservices.

Veriblock-FL | Python, PyTorch, gRPC, Distributed Systems, zk-SNARKs

Apr 2025

- Built fault-tolerant federated learning system with Byzantine-resilient gradient aggregation; zk-SNARK verification prevents gradient poisoning, gRPC cuts inter-node overhead 10% vs REST; maintained liveness under **30% node failure**.

CDataDB | C, B-Tree, Write-Ahead Logging, LRU Buffer Pool

May 2025

- Implemented disk-based relational DB engine in C with B-Tree indexing, WAL crash recovery, and LRU buffer pool manager; supports SQL-subset query parsing — demonstrates low-level systems fundamentals.

TECHNICAL SKILLS

Languages: Python, Java, JavaScript, TypeScript, SQL, C, Go

Backend / APIs: FastAPI, Flask, Spring Boot, REST APIs, GraphQL, gRPC, WebSockets

Distributed Systems: Microservices, Event-Driven Architecture, Kafka, RabbitMQ, Sharding, Rate Limiting

Databases: PostgreSQL, MySQL, MongoDB, Redis, Elasticsearch, FAISS, ChromaDB

Cloud / DevOps: AWS (EC2, S3, Lambda), Docker, Kubernetes, GitLab CI/CD, GitHub Actions, Terraform

Observability: Prometheus, Grafana, Distributed Tracing, MLflow

RESEARCH & PUBLICATIONS

[NeurIPS 2026 — Under Review] SCAR: SLO-Constrained Anytime Reasoning for Language Model Agents — LLM inference controller optimizing accuracy under strict per-query latency SLOs; tested on MATH-500, AIME 2025, GPQA. openreview.net #18900

[Published] Veriblock-FL: A Blockchain-Based Verifiable Federated Learning Framework Using zk-SNARKs — IJSDR Vol. 11 Issue 4, Apr 2026 | Impact Factor: 9.15 | Paper ID: IJSDR2604263

[Preprint] NeuroChainOps — Privacy-Preserving Blockchain MLOps for Federated NAS — sciety.org (ResearchSquare)

ACHIEVEMENTS

6+ production-grade distributed systems deployed with full observability and automated CI/CD on AWS.