AMAR KUMAR

Senior Backend Engineer · Distributed Systems at Scale · AI Integration · Full Lifecycle Ownership

Gwalior, MP, India | +91-9685079691 | amarsk.dev@gmail.com | Linkedin | Github

PROFESSIONAL SUMMARY

Senior Backend Engineer who architects and ships production-grade distributed systems at scale — from HLD to deployment — serving 50K+ concurrent users with 99.9% availability, 40% lower DB load, and 30–40% faster APIs. IBM Certified AI Engineer with a proven track record of integrating LLM-powered microservices into real-world SaaS platforms using Kafka, Kubernetes, gRPC, and event-driven pipelines.

CORE TECHNICAL SKILLS

System Design	HLD · LLD · Microservices · Event-Driven Architecture · Multi-tenancy · Scalability Patterns
Languages	Python · PHP · JavaScript · SQL
Backend	Django, Laravel, REST APIs, RBAC/JWT, gRPC · Kafka (Event Streaming) · Message Queues (Redis Queue, DLQ, Exponential Backoff) · GraphQL · Queues & Workers, Service-Oriented Design
Databases	Postgresql, MySQL (Optimization & Indexing) · Redis (Caching & Pub/Sub) · Schema Design
Cloud & DevOps	AWS (EC2, S3, RDS, CloudWatch) · Docker · Kubernetes (K8s) · Elasticsearch · CI/CD (GitHub Actions)
AI / ML	LLM Integration · AI Pipelines · IBM AI Engineer Certified
Observability	Distributed Tracing · CloudWatch · Performance Profiling · Fault Tolerance
Fundamentals	DSA · SOLID Principles · OOP · Clean Code · Code Reviews

PROFESSIONAL EXPERIENCE

Software Engineer → Senior Software Engineer — Backend & Systems

Sep 2022 – Present
Promoted to Senior SE in 2024

TatvaSoft (Remote)

System Architecture & Scalability

- **Spearheaded HLD and scaling** of a multi-tenant SaaS backend supporting **50K+ concurrent users** with high availability and fault-tolerant design.
- **Engineered high-concurrency REST APIs** handling **~5K concurrent requests** during peak webinar events using Laravel + Redis with zero-downtime performance.
- Achieved **30–40% reduction in API latency** via composite indexing, query refactoring, and intelligent response shaping.
- **Designed a distributed caching layer** (Cache-aside & Write-through with Redis), cutting primary DB load by **40%+** during traffic surges.

Engineering Excellence & AI Integration

- **Developed async event-driven pipelines** using Laravel Queues & Workers for bulk processing, dramatically improving UI responsiveness and system decoupling.
- **Led a major LLD Database Refactoring initiative** — cursor-based pagination and data chunking for **1M+ record datasets**, eliminating OOM errors in production.
- **Integrated Python-based AI microservices** for automated content classification, reducing manual operational overhead by **~60%** via REST-bridged workflows.
- Mentored junior engineers on SOLID principles, clean code practices, and performance profiling, raising overall team code quality.

DevOps & Reliability Engineering

- **Standardized containerized deployments** on AWS (EC2, S3, RDS) with Docker, achieving full environment parity and eliminating configuration drift.

- **Built CI/CD pipelines via GitHub Actions**, reducing deployment lead time and manual error rates significantly.
- **Implemented proactive observability** on AWS CloudWatch — custom dashboards, automated alerts for error rates and queue depths, reducing MTTR measurably.

Software Engineer — Backend

Jul 2021 – Sep 2022

Synram Technolab

Scalability & Database Engineering

- Engineered scalable backend architectures for high-traffic booking platforms, ensuring resilience and availability under peak concurrent load.
- **Boosted database throughput by 25–35%** through advanced indexing, comprehensive query refactoring, and a Redis caching layer for hot datasets.
- Designed LLD components for booking engines — focused on minimizing race conditions and guaranteeing data consistency across distributed transactions.

FinTech & Security Integration

- **Architected resilient payment processing modules** with idempotent transaction handling and automated retry logic, eliminating duplicate billing.
- Implemented secure webhook verification and failure recovery, significantly reducing payment-related support tickets.
- Hardened security posture by implementing MFA, RBAC, and CSP headers; automated input sanitization to mitigate SQL injection and XSS vulnerabilities.

KEY PROJECT — AI-POWERED SMART NOTES & JOURNALING SAAS

Stack: *Django · Django REST Framework · Python · NLP / Transformer Models · PostgreSQL · MySQL · Redis · AWS (EC2, RDS, S3) · Background Jobs · REST APIs · RBAC / JWT*

- **Built a scalable SaaS backend** using Django & DRF, supporting journaling, planners, smart notes, and social features on a modular, clean-architecture foundation with reusable services.
- **Implemented an AI inference pipeline** using transformer-based NLP for note summarization and “chat with notes”, with graceful fallback to rule-based processing, achieving 99.9% system availability.
- **Built a scalable API Gateway** with JWT auth, RBAC, and global rate-limiting to secure premium endpoints and protect against DDoS and resource exhaustion.
- **Designed a subscription state machine** for premium billing — trial-to-paid transitions, webhook-driven sync, invoice management, and automated failure escalation for secure access control.
- **Implemented high-performance Redis caching** with event-driven invalidation, maintaining strict cache–DB consistency; optimised DB queries and background processing for production-scale performance.
- **Engineered distributed task orchestration** — Dead-Letter Queues (DLQ) with exponential backoff for critical email, notification, and async workflows; background jobs for report generation.
- **Designed social & discovery features** — public/private notes, follower/following system, trending notes feed, saved note libraries, and real-time chat & support workflows for user engagement.

EDUCATION

B.Tech + M.Tech (Integrated) — Information Technology

Jul 2014 – Jun 2019

Indian Institute of Information Technology and Management (IIITM), Gwalior

CERTIFICATIONS

IBM AI Engineer Professional Certificate — Machine Learning, Deep Learning & AI Model Deployment Pipelines

KEY STRENGTHS

System Design at Scale · Distributed Systems · AI/LLM Integration · Kafka & Event-Driven Architecture · Kubernetes · Full Lifecycle Engineering · Technical Ownership