

Aman Chadha

+91 8368172857 | aman.chadha.mmi@gmail.com | github.com/ac-mmi

SUMMARY

Backend Engineer with 4+ years of experience building and operating production APIs, transaction systems, and asynchronous backend workflows. Strong background in concurrency, distributed systems, database optimization, caching, observability, and reliable transaction processing. Experienced in designing and debugging production systems with a focus on correctness, performance, and operational reliability.

SKILLS

- **Languages & Runtime:** JavaScript/Node.js, Python, PHP, Go
- **Backend:** Express, FastAPI, REST APIs
- **Databases & Caching:** MySQL, MongoDB, Redis
- **Messaging & Systems:** RabbitMQ, Event-Driven Architecture, Distributed Systems, Concurrency, Transaction Processing
- **Infrastructure & Observability:** Docker, Linux, Grafana, Prometheus
- **Frontend:** React

PROFESSIONAL EXPERIENCE

Supreme Court Bar Association of India

Jun 2022 - Present

Backend Software Engineer

- Built and operated a membership and payments platform serving 29,500+ registered members and 2,800+ monthly active users; unified 4 payment products behind a shared gateway and settlement pipeline that processed INR 1.1+ crore during a peak renewal month.
- Designed idempotent payment settlement and transactional accounting workflows processing 4,649 successful transactions worth INR 1.9+ crore, achieving a 97.7% payment success rate with zero duplicate settlements.
- Architected RabbitMQ-based asynchronous pipelines for payment notifications, Puppeteer-based PDF receipt generation, and OTP delivery, keeping non-critical workloads off the synchronous payment path; handled 5,375+ OTP requests with a 74% verification rate.
- Diagnosed and resolved production race conditions in concurrent transaction processing, ensuring consistent transaction state and sequential receipt generation.
- Optimized read-heavy workflows through caching, database indexing, and query optimization for dues management and reporting.
- Implemented Prometheus and Grafana observability to detect, investigate, and troubleshoot production issues.

Open Source Contributions

llama.cpp - ggml-org <https://github.com/ggml-org/llama.cpp/pull/27891>

- Fixed a performance regression in the KleidiAI inference dispatch path by eliminating repeated mutex-taking initialization from hot-path buffer checks.
- Investigated and reproduced the regression using profiling and benchmarking; improved token generation throughput from ~53 tok/s to ~220 tok/s on Apple M5, while reducing mutex-wait samples from 5,703 to 0.
- Contributed the fix upstream to **ggml-org/llama.cpp**, where the pull request was reviewed, approved, and merged into master.

jQuery.cleanData traversal issue(Bug) <https://github.com/jquery/jquery/pull/5523>

jQuery

- Fixed jQuery event-handling cleanup bug (#5523) and added regression tests to prevent recurrence.

EDUCATION

Delhi University

2019 - 2022

Bachelors of Science, Computer Science (Hons.)

- **GPA:** 8.0
- **Coursework:** Data Structure and Algorithms, Linux, Database Management System, Object Oriented Programming(Java), Computer Networking