

# ABHINAV YADAV

+91 8630087242 Noida, India

[abhi.ay.in@gmail.com](mailto:abhi.ay.in@gmail.com) [linkedin.com/in/abhiyd](https://www.linkedin.com/in/abhiyd) [github.com/abhinav-yadav-official](https://github.com/abhinav-yadav-official)

## EXPERIENCE

---

### Instahyre

*New Delhi, India*

#### SDE III

2024 – Present

- Owned multiple backend domains end-to-end, including opportunity calculation, candidate visibility, recruiter workflows, and background processing pipelines.
- Led large-scale refactors of legacy business logic to improve correctness, performance, and long-term maintainability across core systems.
- Made architecture-level decisions on synchronous vs asynchronous execution, caching boundaries, data ownership, and query placement under production constraints.
- Drove complex production rollouts involving Elasticsearch index changes, schema migrations, data backfills, and feature-flagged cutovers with zero downtime.
- Improved system reliability by identifying and fixing systemic issues across Celery, RabbitMQ, Elasticsearch, and MySQL rather than applying localized patches.
- Introduced automated validation and consistency checks to keep data aligned across APIs, UI, background jobs, and scheduled recomputations.
- Acted as escalation owner for high-impact production incidents and reviewed performance-sensitive, security-critical, and data-integrity-impacting changes across the backend codebase.
- Mentored engineers on debugging production systems, reasoning about race conditions, designing safer migrations, and writing resilient async workflows.

#### Senior Software Engineer (SDE II)

2022 – 2024

- Optimized backend performance across high-traffic systems (Search, Candidates, Jobs, Opportunities).
- Benchmarked Redis, ElastiCache, and Memcached under concurrent load to replace DatabaseCache safely.
- Designed Redis-based caching with TTLs, safe invalidation, and connection reuse for hot APIs.
- Introduced multi-server architecture separating web, Celery workers, and outbound processing to reduce load.
- Improved async throughput by fixing retry storms, queue misconfigurations, and long-running tasks.
- Debugged production issues involving RDS replica lag, deadlocks, CPU spikes, and lock contention.
- Enabled secure production debugging via restricted SSH access, Unix user isolation, and secrets hardening.
- Set up VPN-based internal access and strengthened infra security using AWS Secrets Manager.

#### Software Engineer (SDE I)

2019 – 2022

- Built Employer Team Management with RBAC (Admin/Member roles), auditability, and impersonation safety.
- Fixed multiple authorization bypasses protecting premium features at API and UI layers.
- Improved async reliability by resolving race conditions and non-atomic updates in Django workflows.
- Integrated Datadog and Sentry with Celery for production observability and async error tracking.
- Optimized slow MySQL queries by influencing query planner behavior and index usage.

## TECHNICAL SKILLS

---

<b>Languages</b>	Python, Go, Java, C/C++
<b>Databases</b>	MySQL (InnoDB), Redis, Elasticsearch / OpenSearch
<b>Backend</b>	Django, Celery, RabbitMQ, uWSGI, NGINX
<b>Cloud</b>	AWS (EC2, RDS, S3, CloudFront, Lambda, SNS, SQS)
<b>Concepts</b>	Distributed Systems, Microservices, Caching, Async Systems, Performance Tuning
<b>Observability</b>	Datadog, CloudWatch, Sentry

## EDUCATION

---

B.Tech (Computer Science), AKTU

2015 – 2019