

# ZEBA ALI

Backend Engineer | Java, Spring Boot, Distributed Systems | Oracle DB Internals

Noida, India | zebaali1415@gmail.com | +91-9953098290 | [zebacodes.com](https://zebacodes.com) | [github.com/Zebaali-hub](https://github.com/Zebaali-hub) | [linkedin.com/in/zeba-a-7173251a0](https://linkedin.com/in/zeba-a-7173251a0)

## SUMMARY

Backend Engineer with 3.5 years of distributed systems engineering at Oracle. Architected concurrent execution frameworks at 5 billion rows across Exadata, Exascale, and RAC. Pioneered vector search workloads using HNSW and IVF indexing adopted by multiple engineering teams. Resolved over 10 enterprise production escalations for Fortune 500 clients. Proficient in Java, Spring Boot, PostgreSQL, and system design. Currently building open-source backend infrastructure tools.

## TECHNICAL SKILLS

**Languages:** Java, PL/SQL, SQL, Python, Shell Scripting

**Backend:** Spring Boot 3.x, REST API Design, Microservices, Spring Security, JWT, JPA, Hibernate, Maven, JUnit 5

**Databases:** Oracle Database, Oracle RAC, PostgreSQL, Redis, pgvector, NoSQL

**DB Internals:** MVCC, B-Tree Indexing, Hash Indexing, Connection Pooling, Query Optimization, EXPLAIN Plans, AWR Reports, SQL Monitor, Partitioning, TDE, Data Pump

**AI and Vector Search:** HNSW Indexing, IVF Indexing, Vector Databases, pgvector, RAG Systems, Embedding Search

**Engineering:** Distributed Systems, Concurrency, Multithreading, System Design, CI/CD Pipelines, Agile, Scrum, Git, GitHub, Docker, Linux

## PROFESSIONAL EXPERIENCE

### Member of Technical Staff (MTS, IC2) | Oracle India Pvt. Ltd.

July 2022 to April 2026

Oracle Server Technologies Group | Oracle RDBMS 23ai and 19c | Noida | SMTS (IC3) promotion nominated by manager

- Architected concurrent execution frameworks simulating enterprise workloads at up to **5 billion rows** across Oracle RAC, Exadata, and Exascale, managing **100 simultaneous PDBs** under configurable parallel thread loads. Surfaced deadlock patterns, MVCC edge cases, and kernel failures that influenced Oracle RDBMS 23ai production hardening.
- Sole responder for **over 10 priority enterprise production escalations** from global financial institutions, major technology companies, and pharmaceutical enterprises including Wells Fargo, MERCK, and GE. Analyzed Oracle TRC files, AWR reports, and SQL Monitor traces to classify failures by system layer. Coordinated cross-timezone with US development teams to deliver fixes within SLA commitments.
- Designed and built Oracle team's first parameterized vector search query engine with **over 500 custom HNSW and IVF workloads** benchmarking p50, p95, and p99 latency, recall, and throughput at **5 billion row scale**. Workload adopted as standard library across Oracle RDBMS teams and cross-functional engineering groups.
- Worked with Storch, Oracle's internal Java-based CI/CD automation platform, configuring distributed job orchestration, scheduling execution pipelines across multiple RDBMS configurations, and integrating automated workflows into production release cycles. Platform follows Spring Boot service architecture patterns for job scheduling and distributed execution.
- Worked with ADE (Application Development Environment), Oracle's internal Java-based build and source management platform. Reviewed and validated developer Java source code changes for features owned end-to-end. Configured OLAP GLOBALDB and OLAP SSB Java-enabled schemas and Java runtime environments for analytical query workloads.
- Identified and managed **over 200 critical production defects** across 150 directly filed and 50 lifecycle managed. Shipped 13 fixes to production releases. Owned 4 RDBMS features end-to-end: Non-Positional INSERT, Object Number Reuse, Default VARCHAR2 Sizing, and Nested WITH Clause enabling MongoDB aggregation pipeline translation to Oracle SQL.
- Delivered technical presentations to **over 50 member** engineering organization. Automated data pipeline framework across over 10 enterprise schemas including MERCK, GE, MSTR, HARTFORD, and POSTBANK. Led Agile and Scrum delivery cadence with daily standups and cross-timezone US team coordination. Absorbed departed colleague's full workload without impacting release timelines.

### Software Engineer Intern | Wabtec Corporation

January 2022 to July 2022

- Built Java backend data pipelines and Microsoft SQL Server integrations for rail logistics workflows. Developed Master Data Management automation using Informatica PowerCenter, reducing manual data reconciliation effort by 40 percent.

## PROJECTS

### Log Intelligence Tool | Java 17, Spring Boot 3.2, PostgreSQL, Thymeleaf, Docker

- Open-source Spring Boot service automating log file investigation. Parses Oracle TRC files, AWR reports, SQL Monitor traces, and Java stack traces. Classifies failures by system layer. Generates HTML and PDF root cause reports. Built from 4 years of Fortune 500 escalation analysis. GitHub: [github.com/Zebaali-hub/log-intelligence-tool](https://github.com/Zebaali-hub/log-intelligence-tool)

### Database Stress Testing Framework | Java 17, Spring Boot 3.2, JDBC, PostgreSQL, WebSocket, Docker

- Configurable Spring Boot REST service for stress-testing any JDBC database. Concurrent thread executor with real-time TPS and p50, p95, p99 latency metrics via WebSocket dashboard. PDF and JSON report generation. Productized from 4 years of Oracle concurrent execution engineering at 5 billion row scale. GitHub: [github.com/Zebaali-hub/db-stress-framework](https://github.com/Zebaali-hub/db-stress-framework)

### Vector Search Benchmarking Platform | Java 17, Spring Boot 3.2, pgvector, PostgreSQL, Docker

- Spring Boot REST service benchmarking HNSW versus IVF versus Flat vector search strategies on pgvector. Measures recall, p50, p95, and p99 latency, index build time, and memory across dataset scales. Built on real Oracle HNSW and IVF production knowledge at 5 billion rows. GitHub: [github.com/Zebaali-hub/vector-search-benchmark](https://github.com/Zebaali-hub/vector-search-benchmark)

## EDUCATION

Bachelor of Technology, Computer Science | IGDТУW, Delhi NCR

May 2022