

Dhrumil Bhut

Software Engineer | Backend Systems & Distributed Infrastructure

dhrumilbhut@gmail.com | +91 6358823190 | linkedin.com/in/dhrumilbhut | github.com/dhrumilbhut | dhrumilbhut.com

PROFESSIONAL SUMMARY

Software Engineer with 2.5+ years of experience building distributed backend systems using Node.js, PostgreSQL, Redis, and RabbitMQ. Experienced in designing async processing pipelines, distributed job systems, and reliability-focused architectures that handle concurrency, failures, and scale.

WORK EXPERIENCE

Software Engineer

08/2023 to Present

Zuru Tech India Pvt. Ltd., Ahmedabad, India

- Architected the core backend for an Influencer Management Platform handling content management, analytics, and media distribution using Node.js, PostgreSQL, and SQLite in a production multi-database setup.
- Developed an AWS Lambda API for automated media preview generation, cutting compute costs ~20% and improving media load times ~30%.
- Designed a scheduled YouTube Analytics sync pipeline processing performance data for 1000+ influencer channels, enabling real-time analytics without on-demand API calls.
- Optimized DB schema and query patterns, improving overall application performance by 20% and reducing query execution time by 40%.

PROJECTS

ContextLens (Self-Hosted RAG Observability & Hallucination Debugger)

github | live

- Built a RAG observability stack (FastAPI, Celery, PostgreSQL + pgvector, Redis, Next.js) that decomposes LLM responses into atomic claims, attributes each to its source chunk via cosine similarity, and runs an LLM-as-judge faithfulness check, surfacing per-claim verdicts (faithful, partial, unfaithful, refusal) instead of an opaque pass/fail score.
- Designed a Python SDK with a fire-and-forget context manager: trace logging runs on a background daemon thread, verified at 0.56ms on the happy path and under 1ms when the backend is unreachable or disabled, so instrumentation never adds latency or risk to the caller's request path.
- Calibrated the attribution threshold against an organic 14-query RAG test run (not synthetic data): found the initial 0.75 cosine similarity cutoff misclassified paraphrased claims as retrieval failures, fixed it with a three-band confidence model (low-confidence band 0.65-0.75) and separate refusal detection so LLM declines aren't counted as hallucinations.

PayFlow (Payment Processing System)

github | live

- Engineered an async payment processing system with Node.js, RabbitMQ, PostgreSQL, and Redis: handling retries, idempotency, and concurrency with two-layer locking via Redis SET NX and PostgreSQL FOR UPDATE, ensuring no payment is double-processed even under concurrent worker load.
- Implemented production resilience patterns: circuit breaker for fast-fail on gateway outages, exponential backoff retries via Dead Letter Exchange surviving worker crashes, and a crash-recovery watchdog ensuring zero stuck payments under failure.
- Deployed a multi-service architecture (API + Worker) on Railway with 47 unit and integration tests covering state transitions, race conditions, and webhook conflict resolution.

Influencer Management Platform (Zuru Tech India)

08/2023 to Present

- Engineered a centralized Express.js middleware layer with JWT/session auth, request validation, and dual-layer rate limiting: independently throttling inbound API traffic and outbound YouTube API quota to enforce hard quota boundaries without cross-interference.
- Implemented BullMQ + Redis distributed job scheduling with deduplication, retry logic, and concurrency control; built RabbitMQ consumer pipelines and cron-based schedulers for automated media publishing across Instagram, TikTok, and YouTube, fully decoupling heavy operations from the request lifecycle.
- Built a full asset library with watermark-based vendor distribution: background jobs produce processed copies with 6 placement options, proportional scaling, and edge-anchored positioning across image and video, keeping internal originals unchanged.

SKILLS

Languages

Python | JavaScript | SQL

Backend & API

Node.js | Express.js | REST APIs | AWS Lambda | JWT | Middleware Design | Next.js

Queues & Jobs

RabbitMQ | BullMQ | Redis | Dead Letter Exchange

Databases & Caching

PostgreSQL | SQLite | Redis

Resilience Patterns

Circuit Breaker | Idempotency | Distributed Locking | Exponential Backoff | Concurrency Control

AI / ML Infrastructure

RAG Pipelines | pgvector | Celery | FastAPI | LLM-as-Judge Evaluation | MLflow

Testing

Jest | Unit Testing | Integration Testing

Infrastructure & DevOps

Docker | Docker Compose | Git | GitHub Actions | AWS | Railway

EDUCATION

B.E. Computer Science & Engineering

2020 to 2024

Darshan Institute of Engineering and Technology, Rajkot, India

CGPA: 8.74

CERTIFICATES

IBM AI Engineering, IBM (Coursera)

Machine Learning Specialization, deeplearning.ai (Coursera)