

# Rahul Garg

rgarg2605@gmail.com | linkedin.com/in/rahul-garg-2605 | (+91) 87551-81453

Backend engineer specializing in distributed systems, event-driven architectures, and platform infrastructure at scale.  
6+ years building high-throughput services processing millions of daily events across payments, search, and AI-powered systems.

## Education

---

Indian Institute of Technology, Roorkee – B.Tech in Electrical Engineering

Jul 2016 – Jun 2020

## Skills

---

**Cloud & Infrastructure:** AWS (Lambda, SQS, S3, Kinesis, API Gateway), Kubernetes, Terraform, Docker, Git

**Databases:** PostgreSQL, DynamoDB, Redis, OpenSearch

**Languages:** Python, Java, JavaScript, TypeScript

**Frameworks & Libraries:** Django, React, OpenAI Agent SDK

**AI/ML:** LLM Integration, Agentic Workflows, Multi-Agent Orchestration

## Experience

---

Senior Software Engineer II, Blink Health – Bengaluru, IN

April 2025 – Present

- Architected in-house ticketing and call analytics platform replacing Zendesk and LevelAI, saving \$1M+ annually. Built multi-agent LLM workflows for automated call transcription, quality scoring, and support agent QA across 10K+ daily calls. (SQS, OpenAI Agent SDK)
- Led 3-engineer team to replace legacy IVR with LLM-powered voice assistants, designing end-to-end pipeline with real-time STT/TTS, stateful multi-turn conversations, and ticketing reconciliation. Reduced average call handling time from 8 to 5 minutes and cut escalations by 60% (OpenAI Agent SDK, Sierra).
- Led a team of 4 engineers – conducted code and design reviews, drove cross-functional alignment across product, engineering, and operations, and mentored 2 engineers through successful promotion cycles.

Senior Software Engineer, Blink Health – Bengaluru, IN

Jan 2023 – April 2025

- **Search Infrastructure for Cross-Entity Querying:** Designed centralized Search Service ingesting data from 7+ microservices into OpenSearch via Kinesis and Lambda-managed distributed indexing, replacing database-level aggregation queries. Serves cross-entity searches at 2,000 RPS with p95 latency under 50ms.
- **Central Business Events Platform:** Designed library-based SDK for structured business event emission across 10+ microservices, processing 1M+ events/day into a queryable centralized event store with trace correlation. Reduced cross-service debugging time from hours to minutes by enabling event-level replay and filtering.
- **Payment Orchestration:** Redesigned Payment Service with auth/capture separation, dynamic processor routing, and circuit breaker patterns – processing \$50M+ monthly transactions across 3+ providers at 99.9% payment success rate and reducing transaction failures by 40%.

Software Engineer, Blink Health – Bengaluru, IN

Feb 2022 – Jan 2023

- **Patient Data Migration:** Orchestrated zero-downtime migration of 5M+ patient records across 300+ read paths to a new Patient Store, handling schema translation and consistency guarantees via dual-write rollout, phased backfills, and feature-flag-controlled traffic shifting.

Software Engineer, Amazon India – Bengaluru, IN

Jul 2020 – Feb 2022

- Accelerated invoice approval cycles by 40% across Amazon India's vendor ecosystem by consolidating per-approval alerts into batched daily digests, processing 100K+ emails daily. (AWS Batch, SQS, and Lambda)
- Built end-to-end vendor credentialing portal enabling 100K+ vendors to self-submit compliance documents, covering frontend, backend APIs, and infrastructure provisioning within Amazon's internal tooling stack.
- Migrated 2TB of vendor documents from legacy storage system to S3, reducing storage costs by 60% and improving document retrieval availability from 95% to 99.99%.

Research Intern, Samsung R&D Institute – Bengaluru, IN

May 2019 – Jul 2019

- Built test automation pipeline for mission-critical devices, improving test coverage by 30%. Received pre-placement offer for top performance.