

Govind Sharma

Head of Engineering · Engineering Leadership

govindgovind852@gmail.com · Mumbai, India · govindio.com · linkedin.com/in/govind-io · github.com/govind-io

Open to Head of Engineering and engineering leadership roles · Remote (US/EU overlap) or India

SUMMARY

Engineering leader with **5+ years** across hands-on engineering and running teams. Promoted from founding engineer to Head of Engineering at Stamina.io, leading a **10-engineer team** and owning roadmap, hiring, on-call and cloud budget with **99.9%+ uptime** and weekly releases, after helping grow ARR from **\$250K to \$2M+ in 9 months**. Promoted twice at Loktantra, hiring and leading a **5–7 engineer** team. Ties engineering plans to revenue and customer outcomes, and still writes production code.

EXPERIENCE

Staff Engineer — Head of Engineering · Full-time Contract | [Stamina.io](https://stamina.io) | Jun 2025 – Oct 2026

Promoted: Founding Backend Engineer — Lead (Jun 2025) → Staff Engineer — Head of Engineering (Apr 2026)

AI sales, CRM and marketing SaaS · **1,000+ customers** · **50+ countries** · **\$2M+ ARR** · New York, United States · Remote

- › As Head of Engineering, led the **10-engineer team**, reported to the CEO, and owned the technical roadmap, planning, weekly releases, on-call and incident response; shipped **SSO/OAuth2**, **RBAC** and an **OAuth 2.1 MCP server** for larger accounts.
- › Ran hiring end to end (**20+ interviews**) as the team grew from 4 to 10 engineers; introduced ADRs for cross-team design decisions, architecture standards and code review practices.
- › Kept **99.9%+ uptime** through SLOs, alerting on **OpenTelemetry**, **Prometheus** and **Grafana**, an on-call rotation and blameless postmortems, and made merges safer with a PR check that boots the full 20-service fleet before anything ships.
- › Owned the cloud budget and vendor choices and cut cloud spend by **40%** by scaling queue workers to zero when there was no work; kept LLM spend in check with shared rate budgets across every worker and per-task model routing (**70%** lower cost per lead).
- › Stayed hands-on on revenue systems: **Zara**, the AI SDR agent (**20K leads/day**, doubled reply rates), the workflow engine (**5,000+ workflows/day**) and the credit ledger behind usage billing.
- › As founding engineer (Jun 2025 – Apr 2026), built the core platform while ARR grew from **\$250K to \$2M+** in 9 months, including **Twilio** calling for 200+ customers making 5,000+ calls/day.

Software Engineer II — Engineering Lead | Better Software · Embedded at [NeshAI](https://nesh.ai) | Jul 2024 – Jun 2025

B2B AI SaaS for Fortune 500 sales teams · Remote

- › Led a **5–6 engineer** team embedded at [NeshAI](https://nesh.ai) and the 4–5 person internal platform team; owned delivery, architecture reviews and code quality for enterprise AI products.
- › Delivered the AI product recommendation engine: **LLM-powered retrieval** over **1M+ documents**, p95 retrieval under 800ms, 85% answer accuracy on the eval set, used by **Blue Origin**, **Chevrolet**, **DKSH**, **WRGrace**.
- › Shipped **MCP servers** and AI workbooks rated **5 stars** by customers, and implemented the **SOC 2** controls enterprise clients required.
- › Platform team shipped boilerplates, CI/CD and Terraform templates adopted across all client projects.

Sr. Software Engineer — Tech Lead | Loktantra Mediatech | Jan 2022 – Jun 2024

Promoted twice: Software Engineer (Jan 2022) → Team Lead (Jul 2022) → Sr. Software Engineer — Tech Lead (Jan 2023)

Social and live-streaming platform (Khulke) · Bengaluru, India (Hybrid)

- › Hired and led a **5–7 engineer** team; introduced testing practices and code reviews that cut bugs per sprint from 20+ to 5; wrote the TDDs/RFCs.
- › Built a self-hosted WebRTC SFU on **LiveKit**, a signaling server and a React SDK from scratch, replacing a **\$2,000/month** vendor; scaled to 100,000 concurrent users across the platform at 1,000 req/sec, with rooms of 50+ speakers and 500+ listeners.

- › Cut page load from **4.2s to under 1s** and API p95 from 900ms to 200ms through Redis caching, query tuning, code splitting and optimistic updates.
- › Moved deployments to **Kubernetes** with automated rollouts, cutting deploy downtime from **5 minutes to under 2 seconds**.

Full Stack Developer · ICLEAR CCTV & Security System, Delhi (Remote) · May 2021 – Jan 2022 · Co-built the React frontend of a CCTV livestreaming app; repeat customers rose 30%.

PROJECTS & OPEN SOURCE

- › [FastBowlingLab](#) · Founder, Aug 2026 – Present — AI fast-bowling coach built solo: **MediaPipe** pose estimation on phone video (33 body landmarks, 19 measurements), annotated video and adaptive training plans. **1,200+ real users**, 300+ monthly actives, ~20 paying customers, 2,500+ deliveries analysed. Next.js, React Native (Expo), tRPC, PostgreSQL, Inngest, AWS ECS.
- › [livekit/components-js](#) — Merged contribution to LiveKit's official React components, the library behind the WebRTC platform I built at Loktantra.

SKILLS

Leadership:	Technical roadmap, hiring and interviewing, mentoring , sprint and quarterly planning, cross-functional delivery, stakeholder management, RFC/TDD process, budget and vendor management
Reliability:	Incident management , on-call, SLOs , postmortems, OpenTelemetry, Prometheus, Grafana, Datadog
AI / LLM:	Generative AI, LLMs, agentic AI, AI agents, LangGraph, embeddings, vector search, tool calling, MCP servers, LLM evaluation, guardrails, OpenAI, Anthropic Claude API
Platform:	Node.js, NestJS , TypeScript, Python, PostgreSQL, Redis, Kafka , Google Pub/Sub, AWS (EKS, ECS, Lambda, RDS, SQS/SNS), Kubernetes, Terraform , CI/CD
Security:	OAuth2, SSO, RBAC, SOC 2