

Anindya Chowdhury

Senior Software Engineer

anindya.chowdhury@proton.me | +91 8777439217 | Kolkata, India | linkedin.com/in/anindya-41x3n | github.com/41x3n
anindyac.me

SUMMARY

Senior software engineer with 4+ years building distributed backend systems in Go, TypeScript, and Python. Led partner integration engineering at Videoverse, designing the event-driven, RESTful API and messaging infrastructure that powers 300K+ partner events a day.

PROFESSIONAL EXPERIENCE

Videoverse (AI-powered sports media platform), Senior Software Engineer — Partner Engineering Lead 01/2024 – 07/2026 | Remote, India

- Architected and built the **Partner Layer, an event-driven API platform** giving broadcaster partners self-serve access to stream ingestion, metadata, and webhook-based notifications — replacing a fragile legacy integration, with built-in retry and audit logging for reliability. Cut partner onboarding from **1 week to 1 day at 300K+ events/day**.
- Led a 3-person partner engineering team** covering both inbound broadcaster integrations and outbound publishing integrations. Partnered directly with broadcaster engineering teams to define API contracts and failure-handling strategies, contributing to **\$1M+ in cumulative partner deals**.
- Built an extensible connector framework for third-party publishing integrations, cutting delivery time from **over a month of custom code to about 3 weeks** with full end-to-end testing. The framework let the team ship **20+ integrations** — from generic platforms like OneDrive, Dropbox, and Slack to custom builds for Emerge, Storyteller, Imagen, and Reuters — at a sustained pace of about **2 a month**.
- Designed a **multi-channel data ingestion system** for match, tournament, team, and roster data, supporting custom API integrations, self-serve access through the Partner API, and manual upload depending on the client. Cut per-client onboarding time from several weeks to about a week, at a scale of **200+ tournaments, 2K+ teams, and 30K+ players**.
- Built StreamForge, an event-driven Go microservice that ingests live video for broadcasters whose media asset provider has no playlist export. On each stream's SNS trigger, it spins up an **AWS Fargate container** to pull raw .ts segments and assemble a **real-time HLS playlist**, avoiding a costly custom pipeline on the broadcaster's side. Validated at **50+ concurrent streams**.
- Refactored a clip-variant subsystem — aspect ratio, auto-dubbing, captions — out of an overloaded, few-hundred-field MongoDB collection into normalized PostgreSQL models, with **SNS-based dual-write and sub-100ms replication latency**. Unblocked the team to ship new clip-variant types independently.
- Owned the alerting and monitoring side of a **Datadog-to-Last9/PagerDuty migration** and introduced structured incident management practices — part of a broader effort that **cut infrastructure costs 40% and recurring P0/P1 incidents 80%**.

Let's Dive (AI meeting productivity platform), Software Engineer I (Fullstack) 07/2022 – 01/2024 | Bengaluru, India

- Built a webhook-based calendar synchronization system for Google Calendar and Outlook — reschedules, cancellations, attendee changes, conferencing updates — **across 1,000+ connected calendars**.
- Built HubSpot and ClickUp integrations to auto-publish AI-generated meeting summaries, **adopted by 20+ workspaces**.
- Built an AI companion product end to end, with architecture guidance from senior engineers: users texted a Telegram bot, and an **event-driven pipeline (AWS SQS, Lambda)** routed the message through an LLM via **OpenRouter (Pygmalion-7B)**, converted the response to speech with ElevenLabs, and replied — with Stripe handling subscription billing.
- Built the end-of-year recap feature end to end, including UI and transition animations, and contributed to a broader frontend revamp — a **Google Calendar-inspired dashboard** plus integration-settings and publish-workflow screens.

Let's Dive (AI meeting productivity platform), Software Engineer Intern 08/2021 – 06/2022 | Remote, India

- Built games and interactive apps with Firebase and Svelte, run as iframes inside the main React app, reaching **6,000+ sessions a month**. A shared Svelte component library **cut game development time by nearly 60%**.
- Built the CI/CD pipeline for the games and apps codebase.

EDUCATION

MCA (Computer Application), Brainware University 08/2020 – 07/2022 | Kolkata, India

BCA (Computer Application), Brainware University 08/2017 – 07/2020 | Kolkata, India

SKILLS

Languages & Runtime	Frameworks	Database & Infra	AI / Agents
- Go - TypeScript - Python - JavaScript - Node.js	- NestJS - Express - FastAPI - Django - React	- PostgreSQL - MongoDB - Redis, - AWS, GCP - Docker	- LLMs (Gemini, Claude) - RAG - Google ADK - pgvector - Voyage AI

PROJECTS

Samvad, AI news-intelligence platform (retired) + open-source ingestion engine

- Built and later retired Samvad, a production AI news-intelligence platform (ask.samvadhq.com ⓘ) indexing 26 Indian news outlets for citation-backed semantic search and analysis — sunset after validating the concept.
- Designed a multi-agent retrieval pipeline (Guard, Retriever, Analyst, Critic) that combines hybrid search (pgvector, PostgreSQL full-text, and RRF reranking) with asynchronous report generation, surfacing how different outlets framed or omitted the same story.
- Shipped Samvad as a full production SaaS with authentication and credit-based billing on GCP, and open-sourced the Go ingestion engine ⓘ behind it — a pluggable service that pulls from provider feeds and fans events out to SQS, SNS, Pub/Sub, or webhooks.