

# Ratnadeep Bhattacharyya

Forward Deployed Engineer – AI Systems, Backend Platforms & Distributed Infrastructure

✉ [ratnadeep.bhattacharyya@gmail.com](mailto:ratnadeep.bhattacharyya@gmail.com)

☎ +91 8310823305


🌐 <https://github.com/ratnadeep007>

## Skills


Distributed systems, system design, backend architecture, microservices



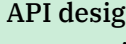
TypeScript, JavaScript, Go, Python (FastAPI, Django)




React, TypeScript, JavaScript, Angular, Flutter, React Native, HTML, CSS, Redux, WebRTC, Twilio, REST APIs, WebSockets




AI agents, LLM applications, prompt engineering, tool calling, RAG pipelines, semantic search, Text-to-SQL, multiagent systems, LLM evaluation (RAGAS), prompt safety, PII filtering



API design, REST APIs, service-toservice communication



Event-driven architecture, asynchronous processing, workflow orchestration



Fault tolerance, system reliability, high availability




Observability, distributed tracing, structured logging, production debugging




Performance optimization, caching strategies



Cloud architecture (AWS: EKS, Lambda, S3; GCP: GKE; Cloudflare: Durable Objects, R2)




Kubernetes, Docker, Terraform




PostgreSQL, MySQL, MongoDB, GBQ, data modeling




Technical leadership, architecture trade-offs, mentoring engineers



Queues, pub/sub systems



Vector search, embeddings



Stakeholder communication, crossfunctional collaboration, requirement gathering



Customer-facing engineering, user workflow analysis, product iteration



Production issue resolution, operational troubleshooting, user feedback-driven improvements



Software Engineer with 7+ years of experience building distributed systems, AI-assisted platforms, and cloud-native backend infrastructure. Strong background in Forward Deployed Engineering-style work, including customer-facing problem solving, rapid feature iteration, workflow customization, production debugging, and end-to-end deployment ownership. Experienced in LLM-powered systems, orchestration workflows, retrieval pipelines, evaluation systems, and asynchronous services using Go, Python, Kubernetes, and AWS. Known for strong system design, fast execution in ambiguous environments, and bridging technical implementation with real-world business requirements.

## Experience

**Springboard (SlideRule Inc.)** India (Remote)  
Software Engineer II (Fullstack) Jan 2025 - Present

- Owned backend services end-to-end, from architecture and development to deployment, monitoring, observability, and production reliability across distributed systems.
- Designed and developed scalable backend services using Go and Python, enabling low-latency and high-throughput workflows across real-time user systems.
- Built AI-assisted evaluation and grading systems using LLM orchestration, contextual retrieval, asynchronous task execution, and scalable backend processing pipelines.
- Built React-based frontend features and internal tools for AI-assisted grading, workflow monitoring, and operational review, improving usability for mentors, students, and operations teams.
- Designed distributed workflow orchestration systems supporting multi-step execution, retry handling, fault tolerance, distributed transaction patterns, and operational traceability across asynchronous services.
- Collaborated directly with mentors, students, operations teams, and non-technical stakeholders to identify workflow gaps, resolve production issues, and improve platform usability and adoption.
- Translated user and operational feedback into backend feature improvements, workflow optimizations, and reliability enhancements across AI-assisted systems.
- Built monitoring and observability tooling using Prometheus, Grafana, structured logging, distributed tracing, and New Relic to improve debugging, alerting, and operational visibility across distributed workflows.
- Improved reliability and workflow execution through resilient service communication patterns, retry strategies, backoff mechanisms, idempotency controls, and fault-tolerant orchestration.
- Optimized backend processing pipelines, reducing grading latency and improving operational consistency by 30%.
- Deployed and operated production workloads on Kubernetes/EKS with CI/CD pipelines, ensuring scalable, observable, and highly available backend environments.

**TarkaLabs** India (Remote)  
Technical Consultant Jan 2024 - Nov 2024

- Built AI-assisted operational workflows for document intake, summarization, classification, and structured data extraction using LLMs and backend orchestration services.
- Designed and implemented Retrieval-Augmented Generation/RAG systems with semantic search and Text-to-SQL capabilities for querying structured and unstructured datasets.
- Developed scalable FastAPI backend services supporting asynchronous execution, orchestration workflows, and external API integrations.
- Built evaluation and validation pipelines using RAGAS with automated scoring, benchmarking, and CI-integrated quality checks for LLM-based workflows.
- Implemented guardrails including PII filtering, prompt safety checks, output validation, and structured monitoring for reliable AI workflow execution.
- Designed multi-step orchestration systems integrating tool execution, contextual retrieval, and workflow automation across distributed backend services.
- Developed an open-source Flutter UI kit at Tarkalabs used to build client applications faster, standardizing reusable components, design patterns, and cross-platform frontend delivery.
- Collaborated with internal teams, external stakeholders, and end users to understand usage patterns, identify workflow gaps, and translate feedback into product improvements across AI-assisted operational workflows.
- Improved observability and operational debugging across asynchronous systems using metrics, logging, tracing, and workflow monitoring practices.

**Nagarro** India (Remote)  
Senior Software Engineer July 2023 - Dec 2023

- Built and optimized backend services for high-volume microservices-based distributed systems supporting 80K+ daily active users, with strong low-latency and reliability requirements.
- Improved system resilience and fault isolation by redesigning service communication patterns, reducing cascading failures across dependent services, and applying Saga-based distributed transaction patterns where appropriate.
- Enhanced backend performance through query optimization, caching strategies, and improvements to asynchronous processing workflows.
- Strengthened observability and production debugging through structured logging, monitoring, and operational visibility improvements.
- Contributed to scalable microservices architecture decisions focused on maintainability, performance, and operational reliability.

**Dentsu World Services** India (Remote)  
Software Engineer July 2021 - June 2023

- Built high-throughput event ingestion systems using Go, processing large-scale distributed data streams for analytics and reporting workflows.
- Designed and implemented event-driven architectures with asynchronous processing pipelines to enable scalable and reliable communication across services.
- Developed backend services on AWS/ECS/EKS, Lambda, and S3 supporting scalable data processing and distributed workloads.
- Built and maintained Airflow-based data pipelines supporting batch and near real-time processing across distributed systems.
- Designed systems handling high-volume ingestion into MongoDB and Snowflake for downstream analytics, reporting, and operational workflows.
- Designed and exposed external-facing APIs for clients to consume analytics data, submit analytics jobs, and integrate reporting workflows into their own systems.
- Built webhook-based notification flows to notify client systems when analytics jobs completed, enabling asynchronous client-side integrations and downstream automation.
- Improved throughput, performance, and cost efficiency through batching strategies, optimized data flows, and caching mechanisms.
- Strengthened reliability and fault tolerance by implementing retry handling, resilient service communication, and failure recovery strategies.
- Improved production debugging and operational visibility across asynchronous workflows using logging, monitoring, and observability tooling.
- Contributed to architectural decisions around backend services, data platform scalability, and distributed system design.

**Cura Technologies Inc** Bangalore, IN  
Software Engineer Dec 2019 - July 2021

- Built backend APIs and distributed services for healthcare systems with strict requirements around reliability, scalability, and data consistency.
- Led migration to Kubernetes/EKS, defining containerization standards, deployment strategies, and scaling practices across backend services.
- Developed a custom Kubernetes operator to automate PostgreSQL provisioning, failover management, and lifecycle operations, improving operational reliability and database resilience.
- Designed and managed AWS infrastructure using Terraform, provisioning scalable and reproducible cloud infrastructure across compute, networking, and storage layers.
- Built backend systems handling high-volume concurrent workloads with strong focus on fault tolerance, low-latency processing, and operational stability.
- Implemented deployment and rollout strategies including rolling deployments and blue/green releases to improve release reliability and minimize downtime.
- Improved backend scalability and query performance through indexing strategies, query optimization, and caching improvements.
- Worked on Flutter-based frontend applications and supported migration from Swift to Flutter to enable cross-platform delivery across mobile and web.
- Built and maintained web-based internal tools to support healthcare operations, debugging workflows, and internal team productivity.
- Collaborated with internal teams, external stakeholders, doctors, nurses, and healthcare operations staff to translate clinical and operational needs into reliable product and platform improvements.

**1024X INNOVATIONS PRIVATE LIMITED** Bangalore, IN  
Founding Engineer (Backend / Platform) Mar 2019 - Oct 2019

- Built backend systems and platform infrastructure from scratch for a 0→1 product, establishing foundational architecture and service design patterns.
- Designed APIs, backend workflows, and service communication models supporting scalable product evolution and distributed workloads.
- Deployed and operated workloads on GKE, defining early containerization, deployment, and scalability practices.
- Handled backend and platform integration for three React Native applications, supporting mobile product development across multiple user-facing workflows.
- Built internal tools to improve operational efficiency, debugging workflows, and internal team productivity.
- Mentored junior developers on backend development, code quality, debugging practices, and delivery execution.
- Collaborated and synchronized with internal teams, external stakeholders, and clients to gather requirements, resolve delivery blockers, and align technical implementation with business needs.
- Owned backend and infrastructure systems end-to-end, contributing to key architectural and operational decisions in an early-stage environment.

**Tapchief (Relevel)** Bangalore, IN  
Product Engineer Oct 2018 - Mar 2019

- Built backend integrations, APIs, and frontend workflows supporting real-time and event-driven user experiences.
- Developed end-to-end product functionality across backend services and Angular-based frontend systems in a fast-paced startup environment.
- Implemented multi-step, resumable onboarding flows in Angular, improving user completion experience across complex registration and setup journeys.
- Built real-time communication features including WebRTC-based video calling and Twilio-powered audio calling for user-to-user interactions.
- Improved deployment stability, debugging workflows, and operational reliability across early-stage platform services.

## Education

**Siliguri Institute of Technology** Bachelor of Technology • 7.79  
Computer Science and Engineering Aug 2013 - Jun 2018

## Projects

**Vidura – Self-hosted Agentic RAG System** Feb 2024 – Oct 2024

- Built a production-grade agentic RAG system with retrieval, memory, and tool execution for contextual AI responses.
- Designed multi-step agent workflows with semantic search, document grounding, and self-hosted deployment support.

<https://vidurabot.com/>

### Agent Orchestration Platform

- Built a full-stack AI agent orchestration platform with Next.js, TypeScript, FastAPI, PostgreSQL, Redis, and Python workers.
- Designed workflow execution, agent configuration, tool execution, Telegram webhook integration, message history, and runtime event mirroring.
- Implemented visual workflow templates, workflow runs, token/cost tracking, and OpenAI-backed agent execution with OpenClaw runtime integration.

<https://github.com/ratnadeep007/agent-orchestration-platform>

### NL-to-SQL

- Built an NL-to-SQL system that converts natural-language analytics questions into safe SQL over structured datasets.
- Developed a Streamlit chat UI, Starlette/Uvicorn API, CLI, runtime schema discovery, SQL validation guardrails, and SQLite-backed demo database.
- Implemented OpenAI-based intent classification, schema-grounded SQL generation, SQL parsing/validation with sqlglot, caching, repair attempts, and automated tests.

<https://github.com/ratnadeep007/nl-to-sql>

### SkycUA – Computer Use Agent

- Built a local computer use agent capable of interacting with system tools through containerized Docker environments.
- Designed execution workflows for task automation, tool orchestration, and controlled local agent operations.

<https://github.com/ratnadeep007/skycua>

### GoClaw – Agent CLI & HTTP Service

- Built an agent system supporting tool execution, LLM interaction, persistent memory, and external integrations.
- Designed CLI and HTTP service interfaces for multi-step reasoning workflows and agent runtime experimentation.

<https://github.com/ratnadeep007/goclaw>