

Debraj Paul

Technical Lead II — Backend & Platform Engineering

+91-94352-21271 · debraj@debrajpaul.com · debrajpaul.com · linkedin.com/in/debraj-paul · github.com/debrajpaul

Highlights: 11+ years building distributed backend systems (media, fintech, food-tech, travel) · ~50% transaction-latency reduction (Swiggy Dineout POS pipelines) · ~95% credit-workflow automation (PT Tujuh Asia P2P lending) · GDPR/CCPA orchestration across 6+ microservices via OneTrust.

Professional Summary

Backend and Platform Technical Lead with 11+ years designing, building, and operating distributed systems across media, fintech, food-tech, and travel. Promoted to Technical Lead II at Warner Bros. Discovery in January 2026. Deep hands-on expertise in Node.js/TypeScript, AWS-native architectures (Lambda, ECS, Glue, OpenSearch), Kafka-based event-driven pipelines, and GraphQL Backend-for-Frontend patterns. Architect of compliance-grade systems: GDPR/CCPA privacy-request orchestration across 6+ microservices, multi-brand async content moderation, and production-aware backfill data pipelines. Strong emphasis on system reliability, API governance, cross-team platform standardisation, and observable, production-ready software.

Core Skills

Languages & Runtime: TypeScript, Node.js (8+ yrs production), JavaScript, Java, Go

Backend & APIs: Express.js, GraphQL (BFF), REST, Spring Boot, Gin — API design and governance at org scale

Cloud (AWS): Lambda, ECS, SQS/SNS (Standard + FIFO + DLQ), S3, DynamoDB, RDS, IAM, Glue, OpenSearch, CloudWatch

Event-Driven & Streaming: Kafka (two-stream contracts, schema evolution), AWS Kinesis, SQS/SNS — at-least-once delivery, idempotency design

Data & Storage: MongoDB / Mongoose, PostgreSQL / MySQL, DynamoDB, Redis, ElasticSearch, OpenSearch, Parquet (Snappy) on S3

Architecture Patterns: Service-registry fan-out · Two-stream Kafka contracts · State-machine async fulfillment · Offline-extract ETL · Cache-aside · Idempotent dedup · Dual-store split (S3 + DynamoDB) · Hot-reloadable per-tenant policy

Infrastructure & CI/CD: AWS CDK, CloudFormation, Docker, Jenkins, GitHub Actions, CircleCI

Compliance & Security: GDPR, CCPA, OneTrust, subject-rights-request orchestration, audit trails, secure data deletion and access controls

Observability: OpenTelemetry, AWS CloudWatch, structured logging, distributed tracing, proactive alerting

Testing: TDD, Jest (9+ yrs), Mocha-Chai, Sinon, Newman, JUnit

Professional Experience

Technical Lead II — Backend & Platform Engineering

Warner Bros. Discovery (via Robosoft Technologies) · Remote (Global) · Aug 2021 – Present

Global content platform · multi-brand (Food Network, CNN, Eurosport, Magnolia, Golf Digest, MotorTrend) · Node.js/TypeScript · Kafka · AWS CDK · GDPR/CCPA-critical

Project: GDPR/CCPA Privacy Request Orchestrator

- Designed and owned a compliance-grade async orchestrator fulfilling GDPR Article 15 (access) and Article 17 (erasure) requests across 6+ PII-holding microservices — architected for regulatory-deadline-bounded delivery, not real-time.
- Implemented a service-registry pattern: the orchestrator reads a registry to determine which downstream services hold PII per user, enabling new services to join compliance by registering — with zero orchestrator changes.
- Applied a two-stream Kafka pattern (request-out + result-in topics) with a subject-request-ID as the idempotency key end-to-end — re-deliveries at any hop are safe.
- State machine tracks per-service status and an aggregate; aggregate flips to complete only when every individual settles, gating the fulfillment callback to the privacy broker. Access exports staged in S3; broker receives a signed URL.

Project: Multi-Brand Async Content Moderation Platform

- Migrated text moderation off a third-party vendor (cost + policy ownership) and absorbed image moderation from a sister team — built and owned both in-house classifiers across a brand-agnostic submission contract.
- Designed a two-stream Kafka architecture (Moderation Request / Response) with per-hop replay isolation — a failure in one classifier never re-drives the full pipeline.
- Implemented hot-reloadable per-brand policy (banned terms, severity thresholds) — trust-and-safety teams ship rule changes without redeploying the moderator. Applied dual-store: image bytes in S3, record metadata + S3 keys in DynamoDB.
- Serving UGC (comments, nicknames, image uploads) from N consumer brand surfaces (web + mobile) with CIAM identity verification at the submission boundary.

Platform-Level Contributions

- Defined and standardised REST and GraphQL API design conventions adopted across all

engineering teams — reducing cross-team integration friction and onboarding time.

- Implemented distributed tracing (OpenTelemetry) across services, improving incident root-cause analysis. Introduced Redis cache-aside strategy reducing DynamoDB read pressure.
- Architected AWS serverless and container infrastructure (Lambda, ECS) using AWS CDK — standardising multi-environment deployments and enabling fault-isolated, cost-optimised scaling.
- Mentored engineers across squads through code reviews, design sessions, and architectural guidance — raising the collective engineering quality bar.

Senior Software Engineer — Payments & Integrations Platform

Swiggy Dineout · Bengaluru, India · Dec 2019 – Jun 2021

High-volume food-tech · Kafka in production · 10+ POS integrations · event-driven pipelines · revenue-critical

- Designed and built a high-performance integration platform connecting 10+ third-party POS systems using Kafka and Redis — reduced transaction latency by ~50% through async event pipelines and optimised caching.
- Built customer segmentation logic on MongoDB and Elasticsearch, enabling behavioural targeting (frequency, spend, activity) — contributed to ~25% conversion lift and ~15% bounce-rate reduction.
- Implemented marketing trigger systems delivering personalised campaigns on real-time segmentation. Designed Kafka topic schemas and a backward-compatible event-evolution strategy.
- Delivered three enterprise-grade AWS microservices contributing to ~25% YoY revenue growth.

Senior Software Engineer — Fintech Platform

PT Teknologi Tujuh Asia · Jakarta, Indonesia · Nov 2018 – Dec 2019

P2P lending · KYC automation · risk decisioning · Node.js/AWS

- Automated ~95% of credit-processing workflows — improved loan-approval accuracy by ~25% and reduced end-to-end processing time by ~50% through system-driven decision pipelines.
- Built OCR + facial-recognition pipelines for KYC and identity verification; developed an NLP-based scoring engine for risk analysis — reduced communication errors by ~90%, improved customer satisfaction by ~20%.

Software Engineer (Backend) — TripBorn · Ekagga Technology · Freelance

Bengaluru, India · Jan 2015 – Oct 2018

- Built booking/reservation backends (TripBorn), enterprise analytics APIs on Java/Spring Boot + PostgreSQL (Ekagga), and end-to-end REST API solutions for multiple clients (Freelance). Full SDLC ownership across all engagements.

Selected Open-Source Projects

Full Budget App — Multi-tenant fintech SaaS budgeting platform on AWS serverless (pnpm monorepo, Node 24, strict TS, AWS CDK). Event-driven ingestion (S3 → SQS → Lambda; DynamoDB Streams → Lambda), hybrid deterministic-NLP + Amazon Bedrock (LLM) categorisation behind cost-aware feature flags, schema-first GraphQL (Apollo Server 5), JWT with refresh-token rotation + reuse detection.
github.com/debrajpaul/full-budget-app · MIT

Certifications

- **AWS Certified Developer – Associate** (2024–2027) — Lambda, IAM, SQS/SNS, API integration, security, cloud best practices
- **Apollo GraphQL Developer – Associate** (2025) — schemas, resolvers, query optimisation, BFF patterns

Education

Master of Computer Applications (MCA) — Tezpur University, Assam, India · 2012 – 2015

Bachelor of Computer Applications (BCA) — Gauhati University, Assam, India · 2009 – 2012