Abhinav Agarwal

aagarwal.gtr@gmail.com · linkedin.com/in/aagarwal-gtr · github.com/aagarwal-gtr

SUMMARY

Tech lead with 9 years in distributed data infrastructure, including 6.5 years on Confluent's Kafka Connect team. Deep expertise in Kafka Connect architecture, Confluent Cloud platform operations, cloud-native Kubernetes deployment, observability at scale, and enterprise disaster recovery. A twice-over founding engineer: built Confluent Connect India from first hire to a 70+ person org, now building Roblox's data engineering function in India from zero.

EXPERIENCE

Roblox — Gurgaon, India

Senior Software Engineer, Tech Lead — ROS-Data — *Mar 2026 – Present*

Founding engineer and tech lead for the ROS-Data team at Roblox's India engineering site.

- Leading AskRaia: a conversational AI agent that surfaces structured insights and visualizations from natural language queries across Roblox's fragmented operational data landscape (Postgres, Grafana, Google Analytics, Glean) — transforming disconnected data origins into a unified intelligence layer for business decision-making.

Confluent — Bangalore, India

Tech Lead, Connect Platform — *Sep 2019 – Mar 2026*

Pod tech lead (∼10 eng) in the 70+ person Connect org.

Disaster Recovery for Managed Connectors — Eng Owner — *2025–26*

- Architected DR system for on-demand connector failover across cloud regions — 55+ enterprise requests, 10+ accounts with >$1M ARR, required for regulated verticals (NIST/DORA). <15 min RTO, 99.95%+ SLAs.
- Drove V1-to-V2 architecture pivot, simplifying offset handling and eliminating DLQ remapping. Full design lifecycle from one-pager through implementation spec and test plan.

Observability Platform Migration — Led entire org — *2025*

- Led stalled Datadog-to-Grafana migration (85+ monitors, 22+ dashboards) to completion before a $6M+ cost-savings deadline. Led full 70+ person org with a core team of 3, running office hours and war rooms.
- Zero observability loss, reduced alert noise across all Connect services.

Control Plane Re-architecture, S2S Security, Config Reliability — *2023–25*

- Migrated Connect APIs from monolith to independent service — zero-downtime cutover via shadow traffic and progressive rollout. 2x dev velocity, 80% fewer API failures, 85% lower p75 provisioning latency.
- Shipped SPIFFE/SPIRE S2S auth across four services (first team at Confluent to complete these flows). Drove org-wide LaunchDarkly reliability: root-caused Segment API outage from flag bloat, set KPIs (80% incident reduction, <30 min MTTR).

Connectors, Public API, Auth, Data Preview — *2019–23*

- Shipped high-throughput connectors (RabbitMQ Sink: 10k+ msg/s, zero major issues over 2 yrs), Connect public API, AuthN/AuthZ service. Led a 4-person feature team delivering a novel connector data preview capability.
- First hire for Connect India. Mentored first intern (hired full-time, promoted to L4), first cohort of 6 new grads, wrote promo packets for 2+ engineers. 250+ interviews. Pioneered AI-assisted on-call: 50–60% of Tier 1 triage automatable; playbooks became standard practice.

Flipkart — Bangalore, India

Software Engineer, Seller Platform — *Jul 2017 – Sep 2019*

- Built and shipped inventory microservices at scale: automated age-driven recalls, capacity orchestration, proxy/auth gateway for Seller API.

EDUCATION

Georgia Tech — M.S. Computer Science, AI specialisation — ongoing

BITS Pilani — B.E. (Hons.) Computer Science — 2013 – 2017

DPS R.K. Puram — AISSCE (Class XII) CBSE — New Delhi, 2013

SKILLS

Kafka: Kafka Connect, connector development, cluster operations, cloud-native deployment (K8s), Confluent Cloud
Data & Infra: gRPC, Spanner, Postgres, Grafana, observability pipelines, disaster recovery **Languages:** Go, Java, Python
Cloud: AWS, GCP, Kubernetes, Docker, Terraform, CI/CD, microservices

TALKS & COMMUNITY

Speaker — Apache Kafka Meetups, Bangalore

Spoke at community meetups on the operational and architectural realities of running Kafka Connect at scale in production — drawing on work done at Confluent.

- *Kafka Connect on Kubernetes* (2025) — architectural patterns for cloud-native connector deployments: operator design, control/data plane separation, self-healing strategies, and upgrade rollout practices.
- *Kafka Connect on Cloud* (2024) — operational practices for managed connector infrastructure: observability pipelines, alert design, reliability patterns, and lessons from running 85+ monitors across a large connector fleet.

Open Source Contributions

Active in the Apache Kafka and data tooling ecosystem, with contributions spanning connector frameworks, analytics, and scientific computing.

- Contributor to Apache Kafka Connect (Confluent fork) — patches and improvements to connector framework internals, informed by 6+ years of deep on-call and production experience running connectors at enterprise scale.
- Past contributor: Metabase (open source BI / data visualization), SciLab (scientific computing), R package BayesBD (Bayesian boundary detection).
- Founding member of the BITS Pilani Firefox Community — organised developer events and drove open source awareness on campus.

INTERNSHIPS & RESEARCH

belong.co — **Bangalore, India**
Intern, Analytics Engineering — *Jul – Dec 2016*

- Owned end-to-end metric automation for key products: data extraction pipelines, visualization, and internal reporting dashboards. Drove adoption of open source BI tool Metabase across the analytics function.
- Redesigned the company's predictive customer health score model and implemented a customer dependency metric from scratch — both used in customer success workflows.

CSIR-CLRI — **Chennai, India**
Intern, Economics Research Division — *May – Jul 2015*

- Built R applications to automate data extraction from research articles and process government import/export datasets for economic analysis. (Guide: N. Vasagam, Sr. Scientist.)
- Developed MATLAB programs for experimental computations and data fitting for polymer sciences applications. (Guide: T. Narasimha, Principal Scientist.)

Mathematical Sciences Foundation — **Delhi, India**
Academic Intern — *Jun – Jul 2011*

- Selected for a national top-50 high school mathematics programme (MSF Delhi / Shiv Nadar University). Studied Ant Colony Optimization algorithms and their applications to graph problems and the Traveling Salesman Problem.

COURSES & CERTIFICATIONS

A sustained interest in machine learning and data science.

- **Machine Learning** — Andrew Ng, Stanford / Coursera. Supervised and unsupervised learning, neural networks, SVMs, and practical ML system design. Completed with distinction.
- **Udacity Machine Learning Nanodegree** — end-to-end ML pipelines, model evaluation, feature engineering, and production deployment patterns. Completed with distinction.
- **Introduction to NLP** — Coursera. Text classification, sequence models, and language understanding fundamentals. Completed with distinction.
- **Data Scientist's Toolbox** — Coursera. Reproducible research, statistical computing, and data pipeline tooling. Completed with distinction.