

Amandeep Srivastava

aman.iitism19@gmail.com | +91 97933 19334 | amandeepsrivastava.in | LinkedIn

SUMMARY

Infrastructure and distributed-systems engineer with **7+ years** building large-scale control planes, data platforms, and developer tooling. Most recently I built a greenfield Infrastructure-as-Code control plane and CLI used by **10,000+ engineers** to manage **~8M resources** across 10+ platforms. I integrate AI into these platforms as a force multiplier on a strong systems foundation, to abstract complexity and deliver measurable efficiency gains.

EDUCATION

IIT (ISM) Dhanbad

B.Tech (Hons.), Computer Science and Engineering

2015 – 2019

CGPA: 9.24 / 10

TECHNICAL SKILLS

Languages: Java, Python, Go, C++

Distributed Systems & Infra: K8s, Crossplane, Helm, gRPC

Data & Streaming: Kafka, Samza, Cassandra, Elasticsearch, MySQL

Graph & Frameworks: Gremlin, SPARQL, JanusGraph, AWS Neptune, Spring Boot, Vert.x

AI Integrations: Claude SDK, MCP, skills, agentic workflows, Cursor

EXPERIENCE

LinkedIn | Senior Software Engineer

Sep 2024 – Present

Core Infrastructure, Control Plane

- Architected LinkedIn's **Infrastructure-as-Code authoring platform** (validation and publishing layers) and a unified CLI, used by **10,000+ engineers** and agentic workflows to manage infrastructure; shift-left validation catches misconfigurations at PR time, cutting apply-time failures by **~90%** across **5,000+ services** (projected **~\$5M/yr** savings).
- Engineered **Long Running Operations** support for the control plane's **Kubernetes/Crossplane** provisioning engine (**~8M resources** across 10+ platforms), modeling multi-step operations as resumable Kubernetes-object state machines; automating these manual flows cut provisioning time **~73%**, extending to **stateful cluster management**.
- Re-architected the manifest model into a **backward-compatible, multi-file structure** that lets a single manifest be split across repositories on demand, enabling file-level resource access control and bounding automation context; rolled it out fleet-wide to **154 services across 90+ teams with zero outages** through an agent-driven migration.
- Built **LinkedIn's first infrastructure-management agent** that converts natural-language requests into validated infrastructure from engineers' dev tools (IDEs, CLIs); a pluggable rules file lets a new platform integrate without building its own agent (**~17 → ~1 day** of dev), while an **autonomous migration agent** onboarded **250+ pre-laC services** onto the control plane with zero outages (**~2 hours → ~20 min** each).
- Developed a **multi-layered AI system** across the platform: an **automated PR-review agent** using GitHub Copilot SDK that reviews every infrastructure change, a **developer on-call agent** that triages across design docs, wikis, Slack, tickets, and code, and user-facing assistant layers in the UI and Slack; the on-call agent matched human judgment on backtests and saved **10–42 minutes per investigation**.

LinkedIn | Software Engineer

May 2022 – Sep 2024

Systems Infrastructure

- Designed and shipped a high-throughput (**~30M QPS**) **resource-insights service** for a distributed NoSQL document store; it became a core dependency of database provisioning (cluster capacity checks), quota changes, deletion (live-traffic safety checks), usage-based quota recommendations, and infra-cost reporting.
- Developed the resource-manager core of a new, centrally managed **infrastructure control plane** designed for pluggable platform onboarding; onboarded the **MySQL platform** as the first tenant, delivering self-serve provisioning with **70%** fewer errors and **53%** lower API latency, and established a pattern now reused across **30+ platforms** (MySQL, Couchbase, Temporal, TiDB).

Goldman Sachs | Software Engineer → Senior Software Engineer

May 2019 – May 2022

Knowledge Management & Research team

- Built a **database-agnostic graph platform** (AWS Neptune, JanusGraph, Gremlin) for fraud detection across firm transaction records, with an **Elasticsearch**-backed search service for autocomplete, fuzzy, and weighted entity lookups over the graph.
- Developed a **Cassandra-backed document store** that cut response times by **~35%** versus the legacy system.

LEADERSHIP & RECOGNITION

Technical leadership: Tech-led a pod of 4 and mentored 8+ engineers; conducted 100+ technical interviews.

Recognition: Google Summer of Code Mentor, DBpedia (2019 & 2020); ranked **1 of 950+** in first-semester examinations at IIT (ISM) Dhanbad.