

Sudeep Gupta

Delhi-(M)+91 8882584824

[in.linkedin.com/in/sudeepgupta90](https://www.linkedin.com/in/sudeepgupta90)
<https://github.com/sudeepgupta90>

sudeepgupta90@gmail.com
sudeep1251@iiitd.ac.in

Summary

Site Reliability and Platform Engineer with 10+ years of experience designing and scaling cloud-native infrastructure, internal developer platforms, and observability systems for large-scale distributed environments, driving reliability, scalability, and engineering efficiency at organizational scale.

Work Experience

Lead Site Reliability Engineer – Cloud Infrastructure Platform

Avalara, Remote(India) | April 2024 – Present

Lead the design and development of internal developer platform and automation tools, while overseeing technical delivery, engineering practices, and scalable solutions that improve developer productivity and operational efficiency across product/engineering teams.

- Architected and developed a metrics-driven SRE compliance platform (Go, Kafka, GitLab, Prometheus) that replaced manual release governance with continuous policy evaluation, reducing deployment lead time by 80% while improving release stability at scale
- Designed and built a Go-based configuration templating engine and validation system for Kubernetes and multi-environment deployments, reducing misconfiguration incidents by 40% and improving deployment hygiene
- Designed AI-driven operational tooling integrating Prometheus metrics, logs, and deployment signals to accelerate root-cause analysis for services deployed on the platform
- Partnered with Platform, Product, and Engineering Leadership to align infrastructure and reliability initiatives with organizational delivery and uptime goals
- Lead a globally distributed team of 4 engineers and increased the team velocity by 30% with Agile coaching, and continuous feedback loops

Technical Stack: Go, Kafka, GitLab CI/CD, GitOps, Kubernetes, Grafana, Prometheus, Agile/Scrum

Senior Infrastructure Engineer

Farfetch, Gurgaon (Delhi NCR) | April 2020 – March 2024

Built and scaled centralized cloud infrastructure and internal platforms supporting Application workloads and Analytics and MLOps workloads, with a focus on reliability, observability, and cost-efficient infrastructure automation for global engineering teams.

- Architected and deployed an Airflow Platform-as-a-Service (Terraform, ArgoCD, Helm) with custom RBAC, secrets backend, centralized logging, and observability; led migration from Google Cloud Composer, reducing platform costs by 70% while improving operational control and reliability
- Designed and implemented a highly available Prometheus observability stack integrated with PagerDuty and DeadMansSnitch, achieving 99.99% platform uptime and saving \$500K annually by retiring Azure Container Insights
- Introduced spot-instance orchestration across Kubernetes and Databricks workloads, optimized GPU and compute utilization, and reduced annual infrastructure spend by \$100K+
- Implemented governance and cost observability frameworks (Databricks Overwatch) to provide automated insights into platform inefficiencies and resource usage patterns

Technical Stack: Azure Cloud(DevOps, Kubernetes, Databricks, Storage, Monitor), ArgoCD, Prometheus, Grafana, Databricks Overwatch, Python, Terraform

Associate, Advanced Data Analytics- Financial Modeling Group

BlackRock, Gurgaon (Delhi NCR) | Jul 2017 – April 2020

Transitioned from hands-on data engineering to building and modernizing large-scale data infrastructure, orchestration platforms, and cloud-native systems supporting Analytics, and MLOps workloads.

- Architected a modular Data Fabric for Equity Research to automate ingestion and storage of multi-source structured and semi-structured datasets, enabling standardized signal generation workflows and reusable compute/analytics layers across research teams (GCP, Python, Flask, MongoDB, Ansible)
- Led migration of on-prem mortgage asset modeling infrastructure to GCP and Airflow (Composer), modernizing legacy batch pipelines, reducing runtime from 48 hours to 10 hours, and significantly improving research iteration cycles
- Designed and implemented a scalable Data Lake platform for low-latency interactive analytics, establishing governance and data organization patterns (Medallion-style layering) to prevent data swamp and support large-scale analytical workloads (Hadoop, Spark, Presto)
- Engineered performance-critical internal tooling (FTPSync) for distributed file system synchronization across HDFS, NFS, and object storage, reducing algorithmic complexity from $O(n^2)$ to $O(n)$ and lowering memory footprint by 20%

Technical Stack: Google Cloud, Airflow, Hadoop, Spark, PrestoDB, Scala, Python, Docker, Kubernetes, MongoDB, Ansible

Senior Data Engineer

Fractal Analytics, Bangalore | Jan 2016 – Jul 2017

Worked on large-scale Big Data and Advanced Analytics systems for strategic enterprise and public-sector clients, focusing on distributed data pipelines, performance optimization, and scalable data infrastructure for analytics-driven decision making.

- Developed distributed financial fraud detection pipelines using Spark, Neo4j, and Python to identify fraud rings and shell entities; optimized pipeline architecture to reduce runtime from 12+ hours to under 2 hours for large graph based datasets
- Designed and implemented Hadoop-based Data Lake and ETL frameworks (Hive, Spark, Ranger, Avro) integrating structured and semi-structured data sources, enabling scalable analytics and warehousing on self-hosted Hortonworks clusters
- Engineered production-grade data processing workflows for high-volume analytical workloads, improving reliability, data consistency, and execution efficiency across client environments

Technical Stack: Apache Spark, Hive, Hadoop, Ranger, Avro, Neo4j, MongoDB, Python, Scala, Hortonworks Data Platform

Earlier Roles Consolidated (Stealth Mode Startup, InnovAccer, Internship)

2014 - 2017

Worked on backend systems, distributed data pipelines, and large-scale data infrastructure across start-up, research, and product environments, building early foundations in automation, cloud-native data workflows, and scalable system design.

- Built distributed data ingestion pipelines, web crawlers, and REST APIs on GCP for large-scale structured data processing
- Developed geo-spatial analytics and backend data platforms using Python, Flask, MongoDB, and MySQL
- Led a small engineering team at a start-up delivering automation-heavy data and ML systems for multiple client domains
- Conducted research on large-scale web crawling and data analysis at IIIT Delhi

Technical Stack: Python, Flask, MongoDB, MySQL, GCP, Distributed Crawlers, REST APIs, Hadoop/Spark

Education

IIIT-Delhi (2012-2014)	: Master of Technology in Computer Science with Major in Data Engineering	GPA 7.3/10
GGSIU, Delhi (2008-2012)	: Bachelor of Technology in Computer Science	79%

Other Achievements

- Kaggle Rank: 8617 (dated: April 2015)
- Ranked in 98th percentile in GATE 2011 conducted by MHRD, Govt. Of India.
- Ranked 1681 (All India Rank) in 7th National Cyber Olympiad, 2007.
- Ranked in 93rd percentile in National IT Aptitude Test, 2009 conducted by NIIT.