

# SATYAJIT ROY CHOWDHURY

Engineering Executive | 20+ years in US & India | Leader in Cloud, AI/ML Infra, and SRE Engineering

San Mateo, 94404 (561) 866-3499 [talk2sroy.ch@gmail.com](mailto:talk2sroy.ch@gmail.com) [LinkedIn](#) [Portfolio](#) [Blogs](#)

## EXECUTIVE SUMMARY

Engineering and Platform leader with 20+ years of experience operating and scaling mission-critical SaaS and AI/ML platforms. Proven at building global SRE organizations, running hyperscale systems serving tens of billions of requests, and improving reliability, security, and cloud efficiency under rapid growth. Trusted partner to product and executive leadership, known for translating technical risk into clear business decisions and developing strong engineering leaders.

## KEY EXECUTIVE ACHIEVEMENTS

- Hyperscale Platform Leadership:** Architected and scaled search and AI platforms processing **30B+** daily requests with **>99.95%** availability, increasing GPU utilization by **40%** through advanced scheduling.
- Growth & Reliability Enablement:** Led platform and SRE strategy enabling **400% traffic growth** and **200%+ SaaS customer expansion** while improving reliability and release velocity.
- Financial & Cost Optimization:** Delivered **22 - 65% cloud cost reductions** across organizations through hybrid architectures and workload optimization, including **30% annual TCO reduction** for large-scale search platforms.
- Organizational & Talent Leadership:** Built and led global engineering organizations of **55+ engineers**, maintaining **<10% attrition** and promoting **30%+** into senior and leadership roles.

## AREAS OF EXPERTISE

### Executive & Organizational Leadership

P&L and Budget Accountability • Global Org Scaling (7 - 55+ engineers) • Manager-of-Managers • Build vs Buy Strategy • Technical Roadmapping • OKRs • FinOps/Cloud TCO Optimization (coralogix)

### Cloud & Platform Engineering

Multi-Cloud (AWS, Azure, GCP) • Kubernetes (EKS, AKS, GKE) • GitOps (ArgoCD) • IaC (Terraform, Crossplane) • eBPF/Service Mesh (Cilium)

### Reliability & Operations

SLO/SLI Design • Incident Command • MTTR Reduction (60 - 73%) • Observability (OpenTelemetry, Prometheus, ELK, Grafana, Jaeger) • Chaos Engineering • Global On-call

### AI/ML & High-Performance Infrastructure

NVIDIA GPU (V100, A100, T4, L4) • HPC Networking (RDMA, InfiniBand, RoCE) • MLOps (Kubeflow, MLflow) • Job Scheduling (Volcano) • Inference Optimization (Inf1, Inf2)

### Product, Delivery & Security

Microservices & API Platforms • CI/CD and Release Engineering • Zero-Trust • FedRAMP High • SOC 2 • PCI-DSS • WAAP/WAF • Data Sovereignty

## TECHNICAL EXPERIENCE

### Director of Engineering and SRE at Arkose Labs - San Mateo, CA

Aug 2024 — Present

Owned platform engineering and SRE strategy for a high-growth fraud detection SaaS, accountable for reliability, security posture, cost efficiency, and scalability as the platform absorbed rapid enterprise-driven traffic growth.

- Amplified platform modernization supporting **7x transaction growth**, re-architecting to EKS-based microservices with **eBPF service mesh** while sustaining **99.9% SLA** under attack traffic.
- Aligned cross-functionally with Product leadership to balance feature velocity and reliability, introducing SLO-based release gates that reduced customer-impacting incidents.

- Directed migration from Cloudflare to CloudFront with Lambda@Edge, reducing **edge latency** by **35%**.
- Served as incident commander for high-severity outages; improved observability and SOC signal quality, reducing **P1 MTTR** by **58%**.
- Scaled the organization from **4 to 16** engineers, managing 2 managers and owning hiring plans and team structure.
- Instituted FinOps governance and tools (coralogix), reducing cloud spend by **22%** while supporting aggressive growth and lowering **effective CPM**.
- Led **build vs buy** evaluations and vendor POCs across cloud, edge, and observability platforms.

## Sr. Director of Product Engineering & Head of SRE at F5 Inc - San Jose, CA Oct 2022 — April 2024

Executive leader with full P&L accountability for Platform Engineering and Global SRE for F5's Distributed Cloud SaaS platform, overseeing reliability, security, and cost efficiency for a globally distributed, security-critical platform.

- Headed engineering and SRE organizations of **55+ engineers** as a manager of managers (3 Directors, 1 Senior Manager), owning hiring plans, headcount allocation, and operating budget across platform and reliability functions.
- Oversaw architectural strategy for a **25+ global PoP** platform delivering multi-cloud networking, WAAP/WAF, and edge services, enabling the platform to absorb **400% growth in attack traffic** while supporting **200%+ customer** expansion.
- Translated Board and C-suite priorities into multi-quarter engineering roadmaps and OKRs, aligning Cloud, SRE, and core product teams to deliver unified releases across the F5 Distributed Cloud portfolio.
- Established governance across infrastructure and platform investments, balancing in-house IP with vendor solutions; slashed annual **TCO** by **30%** while accelerating **feature velocity by 40%**.
- Collaborated with the CISO to deliver FedRAMP High, PCI-DSS, and SOC 2 compliance, implementing policy-as-code, **zero-trust architecture**, mutual TLS, and runtime **eBPF monitoring** to enable expansion into regulated telecom and financial services markets.
- Rebuilt the global SRE operating model into a follow-the-sun structure, reducing on-call burnout and lowering **attrition by 10%** while improving incident response consistency.
- Re-architected observability using ELK, Jaeger, and ML-based anomaly detection, **reducing MTTR** by **73%** and sustaining **>99.92% platform availability**.

## Architect/Technical Leader at Adobe Inc - San Jose, CA

Aug 2018 — Oct 2022

Technical leader for Adobe's Core Search and Sensei Machine Learning platform, owning architecture and infrastructure for hyperscale ML/AI and search workloads.

- Led a team of 12 senior platform and infrastructure engineers, setting architectural direction for Adobe's Core Search and ML platforms used across the Search and Sensei ecosystem.
- Directed the build of a **hybrid GPU/CPU infrastructure** spanning AWS, Azure, and on-prem environments, reliably serving **~30B daily API requests** with **>99.98% availability** for search and sensei platform.
- Designed and optimized large-scale **NVIDIA V100/A100 HPC** clusters with **RDMA InfiniBand, MIG**, and GPU-aware scheduling; integrated MLOps (Kubeflow, Volcano), increasing cluster **utilization by 38%**.
- Led migration from managed AWS Elasticsearch to a self-managed hybrid architecture (18 clusters, **10B+ documents, 6K/sec ingestion**), delivering **30% licensing cost** savings while improving operational control.
- Architected a **unified multi-cloud Kubernetes platform**, consolidating fragmented environments into 15 multi-tenant clusters and eliminating **90% of cluster sprawl** using **Cilium** (eBPF); the reference architecture was later adopted across Adobe.
- Enforced data lifecycle and tiering policies across search and ML pipelines, lowered cloud **storage costs by 65%** while maintaining **sub-5ms P95** latency for latency-sensitive search workloads.

## Tech Leader at Macys.com - San Francisco, CA

Nov 2016 — Aug 2018

Technical leader for enterprise-wide CI/CD and platform modernization supporting Macy's and Bloomingdale's e-commerce platforms.

- Led a team of 15 engineers to design and deliver a **company-wide CI/CD** platform adopted across Macy's and Bloomingdale's engineering organizations.
- Designed a modern deployment platform using Jenkins pipelines, Spinnaker, Kubernetes and in-house orchestration tools, enabling **100+** tested infrastructure deployments per week and reducing deployment time from **2 - 3 days** to under **1 hour**.
- Implemented blue-green and canary deployment strategies, enabling near-zero-downtime releases for revenue-critical e-commerce workloads.
- Designed a hybrid cloud architecture spanning GCP (GKE), AWS (ECS/EKS) subsequently, and on-prem VMware Tanzu, enabling consistent workload execution across environments.
- Introduced early **GitOps** workflows using Flux CD, improving deployment consistency and reducing infrastructure provisioning errors during peak retail traffic.

---

## EARLIER EXPERIENCE

---

- **Prior Engineering Roles (2009 – 2016):** Workday, Chegg, RocketFuel, Adobe, Saba Software, CMC
- **International Engineering Roles (2004 – 2009):** Autonomy, SDG Group, HCL Tech, vCustomer (India)

---

## EDUCATION

---

- **Bachelor of Science** | Mahatma Gandhi Kashi Vidyapith | Varanasi, UP, India | June 2002
- **Leading Effective Decision-Making (certification)** | Yale School of Management | Online | October 2021