

ANAND KUMAR

San Jose, CA • 408-364-6990 • anand.kumar48@gmail.com • [LinkedIn](#) • [GitHub](#) • [Website](#)

Engineering Manager, Infrastructure

Hands-on Engineering Manager with 18 years operating large-scale, cloud-native, containerized platforms and data systems. Led infrastructure teams supporting nine-figure revenue lines and big-data fleets at internet scale — across gaming, AdTech, and AI/ML workloads. Deep, hands-on Kubernetes, Terraform, AWS, and CI/CD depth. Most recently founded the infrastructure stack at an AI-gaming startup.

SKILLS

- **Containers & Orchestration:** Kubernetes, EKS, Helm, Docker, Rancher, Istio
- **Infrastructure as Code:** Terraform, Terragrunt, Chef, Puppet, GitOps
- **Cloud (AWS):** EKS, ECS, EC2, ALB, API Gateway, RDS, Route 53, Lambda, Cloud Map, Amplify, Secrets Manager, IAM, CloudWatch
- **CI/CD:** GitHub Actions (self-hosted runners), GitLab CI, ArgoCD, Jenkins
- **Observability:** Splunk, OpenTelemetry, CloudWatch, Grafana/Loki, AWS X-Ray
- **Distributed Data:** Hadoop, Spark, Solr, YARN, Kafka, Hive, HBase, ZooKeeper
- **AI/ML Infrastructure:** Together AI inference, multi-tier model routing, GPU workload patterns
- **Languages:** Python, Shell

EXPERIENCE

Infrastructure Manager — ErthAI

Oct 2022 – Present

San Jose, CA

- Founding infrastructure hire at a Mark Pincus AI-gaming company. Recruited and built a 5-member infrastructure team, shipping two pre-PMF products: a multi-platform Discord poker game (peak 30K DAU) and Stem Studio, an LLM-assisted no-code game engine.
- Owned the AWS footprint end-to-end with Terraform + Terragrunt: EKS, RDS, Route 53, Lambda, Amplify previews, ALB Ingress.
- Stood up the SRE practice: defined SLIs/SLOs per service, instrumented BetterStack for synthetic uptime monitoring and a public status page, ran a follow-the-sun on-call rotation with time-based paging escalation, and led blameless post-mortems for Sev1 incidents.
- Re-architected compute, storage, and ingress (including aggressive cross-AZ traffic reduction) to deliver a 70% reduction in monthly AWS spend.
- Built the GitOps delivery platform powering multiple codebases: ArgoCD, GitHub Actions, Sealed Secrets, and per-PR ephemeral environments.
- Cut CI/CD build time from 4–5 minutes to 30–45 seconds (~85% reduction), unblocking engineering iteration speed.

Sr. Principal DevOps Engineer — Yahoo

Jun 2016 – Oct 2022

Sunnyvale, CA

- Built and led a team of 7–10 engineers (US + LATAM) inside Yahoo's AI & Data organization, operating the Orion data platform across multiple LATAM countries to support an AdTech business generating ~\$150M in annual revenue.
- Architected and operated a Hadoop ecosystem processing 100B+ records daily across 300+ nodes (Spark as the primary compute engine, Solr for search, plus YARN, Hive, Kafka, and HBase), powering workloads like Mobile Marketing Insights for audience segmentation.
- Led the RedHat–Yahoo integration, starting with 5G Mobile Edge Compute (MEC): deployed enterprise edge platforms with AI/ML capabilities to telco customers.
- Owned platform reliability, capacity planning, and on-call practices across a distributed organization spanning Sunnyvale and multiple LATAM sites.

Senior DevOps Engineer, Big Data — Zynga

Apr 2011 – Jun 2016

San Francisco, CA

- Architected and operated a multi-tenant Splunk platform spanning bare-metal and AWS, indexing 15+ TB/day from 50,000+ forwarders with high resilience.
- Built and operated multiple HA Hadoop, HBase, and ML clusters totaling ~1 PB of storage, powering game telemetry, analytics, and ML across Zynga's portfolio.
- Designed multi-tenant isolation that let multiple game studios share the platform while preserving data segregation and per-tenant performance.

DevOps Engineer — Dell SonicWALL

Jan 2008 – Apr 2011

San Jose, CA

- Managed the engineering datacenter and built/operated observability platforms (Nagios, Cacti) for the engineering organization.
- Stood up virtualization infrastructure that supported product test environments across the engineering team.

OPEN SOURCE

Stem Studio Engine (github.com/Stem-Studio/Engine) — open-source AI-assisted, no-code web game engine.

- TypeScript/WebGPU core. Ships complete games through natural-language prompts; the open-sourced platform behind ErthAI's Stem Studio.

INDEPENDENT PROJECTS

PicCanvas (trypiccanvas.com) — solo-built AI image-generation product

- Production inference on Together AI with a multi-tier model-routing layer that dispatches requests across model sizes to balance cost, latency, and quality.

EDUCATION

M.S., Electrical & Computer Engineering — Illinois Institute of Technology

Aug 2006 – Dec 2008

Chicago