SANTHOSHKUMAR J

Platform Engineer | Cloud Engineer | Site Reliability Engineer (SRE)

Location: Chennai, India | **Phone:** +919865295334 | **Email:** santhoshj.off@gmail.com |
LinkedIn: linkedin.com/in/sk-santhosh-j | **Website:** sk-santhosh.info | **GitHub:** github.com/sk-santhosh

SUMMARY

Platform Engineer, Cloud Engineer, and Site Reliability Engineer (SRE) with 14+ years of experience designing, building, and operating cloud-native platforms on AWS, Microsoft Azure, and Google Cloud Platform (GCP). Proven record of leading 30+ cloud transformation and modernisation projects across multiple industries, engineering platforms serving more than 10 million users at 99.9% availability, and reducing cloud infrastructure costs by up to 35% through FinOps governance. CKA certified, with Go engineering experience including a custom Kubernetes operator built with controller-runtime. Led and mentored cross-functional teams of 20+, spanning Platform Engineering, SRE, DevOps, and Software Development. Delivered large-scale Kubernetes modernisation programmes, cloud-to-cloud migrations of 3 TB+ of production assets, and enterprise observability platforms spanning metrics, logs, traces, and APM across 20+ services and 15+ engineering teams, with CI/CD pipelines reducing deployment effort by 80% and provisioning effort by 70% through IaC and GitOps adoption. Published practitioner-level technical writing at sk-santhosh.info. Open to relocation to Sweden; visa sponsorship required.

EXPERIENCE

Co-Founder & Platform Engineer
Jan 2022 – Present

SentinelFox Pvt Ltd
Chennai, India

- Joined as the Founding Platform Engineer, leading the platform's technical strategy and end-to-end execution while successfully delivering 30+ customer projects. Architected cloud infrastructure, validated product concepts through PoCs, led and mentored 20+ engineers, drove Agile delivery (sprint planning, backlog refinement, daily stand-ups, sprint reviews, and retrospectives), conducted architecture and code reviews, authored technical and operational documentation, and established reusable Terraform modules, Kubernetes platforms, CI/CD templates, and clean, modular, maintainable Infrastructure as Code that reduced operational complexity and accelerated developer productivity.
- Engineered multi-cloud platforms across AWS, Microsoft Azure, and Google Cloud Platform (GCP) supporting workloads serving more than 10 million users at 99.9% availability.
- Deployed and operated a fully self-hosted LGTM observability stack (Loki, Grafana, Tempo, Mimir) on Kubernetes — ingesting 20 GB/day of logs and 10 million active metric series across 7+ customer environments, with 90-day retention on object storage.
- Built GitOps-driven deployment pipelines using Argo CD, GitHub Actions, Jenkins, and Azure DevOps, reducing manual deployment effort by 80%.
- Implemented Infrastructure as Code (IaC) standards with Terraform, Kustomize, and AWS CloudFormation, reducing provisioning effort by more than 70% within 2 months of rollout — eliminating configuration drift and enabling reproducible environments.
- Led FinOps governance initiatives reducing cloud infrastructure costs by up to 35% through rightsizing, autoscaling, Reserved Instance planning, Karpenter-based Spot optimisation, and tagging policies.
- Reduced CI costs by 60% by migrating from GitHub-hosted to self-hosted runners with Karpenter-provisioned Spot capacity, cutting average build times from 12 to 6 minutes through persistent dependency and Docker layer caching.
- Established incident management processes covering on-call rotations, blameless postmortem culture, root-cause analysis, and error budget tracking across 15+ engineering teams — reducing MTTR across 10+ customer environments.
- Defined and implemented Service Level Objectives (SLOs) for critical services, established error budget tracking, and led per-customer SLO evolution — proactively reducing SLO burn through dashboard design, alert tuning, and noise reduction. Published implementation guide at sk-santhosh.info/blogs/implementing-slos-and-error-budgets.
- Managed and scaled Kubernetes clusters (GKE, EKS, AKS) with Karpenter, Helm, multi-zone deployments, and canary rollouts (CKA certified) — owning the full lifecycle from provisioning to decommission.
- Designed and implemented Zero Trust workload identity on Kubernetes — enabling automatic mTLS between services using X.509 certificates with no shared secrets, node attestation, and automated certificate rotation. Published deep-dive at sk-santhosh.info/blogs/zero-trust-workload-identity-spiffe-spire.
- Built a Kubernetes operator in Go using controller-runtime to automate workload identity registration — reconciling Deployments against an identity gRPC API with finalizer-based cleanup, RBAC markers, custom Prometheus metrics, and envtest coverage. Published at sk-santhosh.info/blogs/kubernetes-operator-spiffe-workload-registration.
- Designed and implemented AWS multi-account cloud architectures using landing zone principles, enabling secure workload isolation, centralized governance, and scalable cloud operations.
- Implemented AWS Control Tower, organizational guardrails, and Service Control Policies (SCPs) to enforce security, compliance, and operational standards across AWS accounts.
- Architected enterprise AWS networking using Amazon VPC, Transit Gateway, PrivateLink, Site-to-Site VPN, Direct Connect, NAT Gateways, and hybrid connectivity patterns.
- Led DevSecOps programmes covering vulnerability remediation, compliance hardening, and cloud security governance — resolving 50+ CVEs across 10+ customer environments and establishing security baselines aligned with CIS benchmarks.
- Executed a cloud-to-cloud migration of 3 TB+ of application assets with minimal downtime and no major service disruption.
- Migrated legacy virtual-machine (VM) workloads to Kubernetes and designed highly available platforms using Horizontal Pod Autoscaling, multi-zone deployments, canary rollouts, and automated recovery.

- Deployed self-hosted GitHub Actions runners on Kubernetes using Actions Runner Controller (ARC) — autoscaling ephemeral runner pods from 2 to 10 based on queued workflow demand, executing 1,000+ CI/CD jobs per month.
- Built and led a cross-functional team of 20+ engineers spanning Platform Engineering, Site Reliability Engineering (SRE), Cloud Engineering, and Software Development — mentoring engineers and driving knowledge sharing across global communities.
- Administered Linux-based cloud infrastructure (Ubuntu, Amazon Linux, Debian) across AWS, GCP, and Azure — including performance analysis, system-level troubleshooting, capacity planning, and environment lifecycle management.

Senior Software Engineer
Nov 2017 – Dec 2021

Nextwave Multimedia Pvt Ltd
Chennai, India

- Engineered backend systems and distributed platforms supporting more than 10 million active users at 99.9% availability.
- Designed and developed multiplayer gaming backends supporting up to 10K concurrent players with real-time communication and state synchronisation.
- Re-engineered the tournament platform (live scoring, rankings, and leaderboards) to launch and configure tournaments directly from the backend with artifacts served via S3/CDN — eliminating the per-tournament app release previously required for each new tournament and enabling many public and private (invite-only) tournaments to run in parallel.
- Led migration of legacy PHP applications (200K+ lines of code) to a Laravel-based architecture, improving response times by 60% — reducing average API response time from ~800ms to ~320ms.
- Implemented continuous integration and continuous delivery (CI/CD) pipelines, reducing build and release time from 1hr to 5min.
- Improved application performance through MySQL query optimisation and database architecture enhancements, reducing average query time by 90%.
- Established reusable engineering frameworks and coding standards adopted by a team of 5+ developers, improving maintainability and developer productivity.
- Partnered with engineering leadership on architecture reviews, scalability planning, and technical roadmaps.

Full Stack Developer
Jan 2017 – Oct 2017

eBuilders
Chennai, India

- Developed enterprise web applications using PHP, Node.js, JavaScript, and relational databases.
- Designed and implemented REST APIs and backend services.
- Built Internet of Things (IoT) and automation solutions using Arduino and embedded hardware integrations.
- Contributed across engineering, development, testing, and deployment activities.

Embedded Systems Engineer
Jun 2013 – Dec 2016

Simple Labs
Chennai, India

- Designed and developed embedded hardware and firmware for industrial and commercial applications.
- Worked across the full product lifecycle, from design and development to testing and deployment.
- Built microcontroller-based systems integrating hardware and software components.
- Supported production deployments and troubleshooting activities.

SKILLS

Platform Engineering & SRE	Platform Engineering, Site Reliability Engineering (SRE), High Availability, Disaster Recovery, Incident Management, Capacity Planning, Toil Reduction, Service Level Indicators / Objectives (SLI/SLO), Reliability Engineering, Error Budgets
Cloud Platforms	AWS (EC2, Auto Scaling Groups, ECS, Fargate, Lambda, RDS, CloudFront, Organizations, Control Tower, Well-Architected Framework, Cost Explorer, Budgets, Savings Plans), Google Cloud Platform (GCP), Microsoft Azure (Virtual Machines, Virtual Machine Scale Sets), Cloudflare, Multi-Cloud Architecture, Landing Zone Architecture, Multi-Account Governance
Containers & Orchestration	Kubernetes (GKE, EKS, AKS), Docker, Helm, Karpenter, Kustomize, Kubernetes Operators, controller-runtime, Pod Security Standards, MinIO, Temporal.io
Infrastructure as Code	Terraform, Pulumi, AWS CloudFormation, Kustomize, Infrastructure as Code (IaC), Automation Engineering, Ansible (familiar)
CI/CD & GitOps	GitHub Actions, Jenkins, Azure DevOps, Argo CD, GitOps, FluxCD, Self-hosted GitHub Actions runners (ARC), Deployment Automation
Observability & LGTM Stack	Prometheus, Grafana, Loki, Tempo, Mimir, OpenTelemetry, OTel Collector, Fluent Bit, SigNoz, New Relic, Amazon CloudWatch, Azure Monitor, Google Cloud Operations Suite, ClickHouse
Security & Zero Trust	Zero Trust Architecture, SPIFFE/SPIRE, mTLS, X.509 certificates, IAM (AWS IAM, GCP Workload Identity, Azure AD), Kubernetes RBAC, Network Policies, DevSecOps, CVE remediation, CIS benchmarks
Networking & Traffic Management	NGINX Ingress Controller, Kubernetes Gateway API, TLS/SSL, VPC Networking, Transit Gateway, PrivateLink, Direct Connect, Internal Load Balancing, DNS, Bastion Architecture
Databases	PostgreSQL, MySQL, MongoDB, Redis, Elasticsearch, Neo4j, ClickHouse
Programming Languages	Go (controller-runtime, gRPC, go-spiffe/v2), Python (working knowledge), JavaScript, TypeScript, Node.js, Shell Scripting
Ways of Working	Agile (Scrum), Sprint Planning, Backlog Refinement, Architecture Reviews, Cross-functional Collaboration, Technical Mentoring

EDUCATION

Diploma in Mechatronics (Incomplete)
Nettur Technical Training Foundation (NTTF) · Vellore, India · 2010 – 2013

CERTIFICATIONS

- **Certified Kubernetes Administrator (CKA)** — The Linux Foundation · Dec 2022
- **Google Cloud Professional Data Engineer** — Google Cloud · Dec 2024

PROJECTS

- From Branch to Production: Automating Multi-Environment Deployments with GitHub Actions and Argo CD on EKS
- Reducing Kubernetes Costs by 35% with Karpenter and Spot Instances
- Implementing SLOs and Error Budgets in Production
- Zero-downtime Kubernetes Upgrades: A Production-ready Strategy
- More articles

LANGUAGES

English	Professional working proficiency
Tamil	Native proficiency