

Somasundar Venkatesh

AI-Augmented Site Reliability Engineer · Platform Automation · GenAI Systems
(469) 534-2089 | somasundar.v@gmail.com | [linkedin.com/in/somasundarv](https://www.linkedin.com/in/somasundarv) | Bay Area, CA

SUMMARY

Senior SRE with 17+ years of experience building and operating large-scale, mission-critical infrastructure at Apple and Netskope. Deeply embedded practitioner of Generative AI in SRE workflows — having designed and shipped agentic AI systems that autonomously perform RCA, on-call triage, monitoring gap analysis, and infrastructure troubleshooting. Reduced MTTD by 75% through AI-driven platform automation. Proven ability to bridge traditional reliability engineering with modern AI tooling to dramatically accelerate incident response, reduce toil, and improve platform observability.

CORE COMPETENCIES

Agentic AI Systems · GenAI Workflow Automation · Platform Reliability · SLO/SLI & Error Budgeting · Incident Management · Root Cause Analysis · Observability & Monitoring · On-Call Automation · Production Readiness · Capacity Planning · Kubernetes & Container Ops · Toil Reduction & Automation · CI/CD Pipeline Engineering · Cloud Infrastructure (AWS) · Distributed Systems · Security Automation

EXPERIENCE

Senior Site Reliability Engineer

Feb 2022 – Present

Netskope · Santa Clara, CA

AI & Automation Initiatives

- Designed and built an Agentic AI RCA Generator that autonomously scans Slack incident channels, queries live Grafana dashboards via MCP, reads incident tickets and monitoring snapshots, and produces structured, template-compliant root cause analysis reports — eliminating hours of manual post-incident documentation
- Developed an Agentic AI Platform Troubleshooting System that simultaneously queries live signals across 13 microservices; reduced Mean Time to Detect (MTTD) by 75% across P1/IMF incidents for infrastructure and network-layer issues
- Built an AI-powered On-Call Triage Bot processing ~45+ Opsgenie alerts/month, triaging by severity and blast radius, cross-referencing runbooks in Confluence, and proposing remediation steps — with a self-improving feedback loop that tightens accuracy over time
- Developed a suite of Claude Code skills automating high-toil SRE tasks including capacity analysis, security vulnerability scanning, ticket validation and automated Jira updates, and Grafana dashboard generation
- Built an AI-powered Monitoring Gap Analysis Tool that ingests service configurations and existing alert definitions, identifies blind spots in observability coverage, and auto-generates recommendations and dashboard scaffolding

Platform Reliability & Infrastructure

- Owned end-to-end reliability across 3 teams — including the Governance & Compliance stack — on Netskope's cloud-native SASE platform serving 12M+ enterprise customers, maintaining uptime SLAs across distributed microservices
- Onboarded monitoring for each team — building dashboards, alerting, and uptime tracking — and created reusable, generic dashboard templates adoptable by any team with minimal changes, standardizing observability and accelerating onboarding
- Defined SLOs/SLIs and managed error budgets to drive system reliability, using burn rates to balance reliability investment against feature velocity
- Covered on-call across multiple teams, owning major production incidents end to end — debugging and root-causing issues, then driving blameless postmortems with systemic fixes and action items tracked to closure
- Performed production-readiness reviews ahead of releases to surface reliability risks before deployment
- Drove capacity planning and applied Kubernetes configuration best practices, right-sizing pod resource requests/limits for cluster stability and resource efficiency
- Led an end-to-end API migration from one service to another — evaluating feasibility and scaling pods through capacity analysis to absorb the additional traffic
- Reduced operational toil by building a suite of automation tools across recurring SRE workflows
- Built a metrics pipeline sourcing from the company Datalake (AWS Athena, QuickSight) feeding dashboards used by management and leadership to drive business decisions
- Managed load balancing (HAProxy) — updating and fine-tuning configurations to reduce service impact and minimize customer-facing issues
- Partnered with the platform team to upgrade load-balancing machines and software to the latest versions, minimizing customer impact

- Surfaced numerous code bugs through code review, preventing defects from reaching production
- Maintained and regularly updated runbooks and operational documentation to keep on-call response fast and consistent
- Worked within a Scrum methodology — grooming, triaging, and performing due diligence on tickets in the team queue

Site Reliability Engineer / Technology Lead

Nov 2015 – Feb 2022

Apple · Sunnyvale, CA

- Ensured 24x7x365 stability of enterprise security infrastructure underpinning Apple's global product ecosystem, including the JMETS (Joint Mobile Engineering Technology) platform
- Mentored and guided 9 engineers on reliability practices, SLI/SLA design, and deployment safety standards; served as primary on-call escalation point for enterprise security infrastructure incidents
- Architected and maintained automated deployment pipelines with monitoring, failover, and self-healing capabilities across heterogeneous server and security appliance environments
- Built and maintained CI/CD pipelines using Jenkins, Spinnaker, and Ansible enabling zero-downtime deployments and automated rollback
- Partnered with software engineering teams to define deployment lifecycle standards, SLIs/SLAs, and observability baselines; built dashboards and alerting frameworks to enforce them
- Led and coordinated cross-timezone engineering teams on reliability initiatives spanning multiple product lines and internal platforms
- Delivered root-cause analysis and production support resolution for critical incidents, consistently meeting targets

Technology Analyst

Dec 2012 – Oct 2015

Apple · Mangalore, India

- Delivered full-lifecycle software projects covering requirements, design, development, testing, and production support
- Onboarded a new enterprise system and drove post-production issue resolution, reducing critical defect backlog by 40%

Software Engineer

Oct 2011 – Dec 2012

Turning Point Global Software Solutions · Bangalore, India

- Developed core modules for a Telecom Lifecycle Management System handling inventory reconciliation for enterprise carrier customers

Software Engineer

May 2009 – Oct 2011

Accenture · Hyderabad, India

- Built inbound/outbound file processing pipelines for Bank of America's Singapore File Request Processing project, converting multi-format inputs to GBS-compatible output

TECHNICAL SKILLS

AI & Automation: Claude Code, Agentic AI Systems, GenAI Workflow Design, LLM-powered Tooling, Prompt Engineering

SRE & Observability: Grafana, Prometheus, PromQL, Alertmanager, VictoriaMetrics, Sumo Logic, Splunk, PagerDuty, Opsgenie, SLO/SLI & Error Budget Design, Incident Management, Postmortem Facilitation

Infrastructure: AWS, Linux, Kubernetes, HAProxy, Clickhouse, Mongo, MariaDB, Kafka, Cassandra, CouchBase, Ansible, Jenkins, Spinnaker

Data & Analytics: Datalake, AWS Athena, AWS QuickSight

Languages: Bash/Unix, Python, Java, SQL

DevOps & VCS: Git, Scrum, IntelliJ, Eclipse, VSCode

EDUCATION

Master of Computer Applications (MCA)

2006 – 2009

Sri Venkateswara University · Tirupati, India

Bachelor of Science — Mathematics, Physics & Chemistry

2003 – 2006

Sri Venkateswara University · Tirupati, India