

SAROJ NAYAK

Site Reliability Engineer | Forward Deployed Engineer | AIOps · MLOps · Agentic AI

Bangalore, India • +91 8884311573 • realsarojnayak@gmail.com • linkedin.com/in/saroj810 • github.com/sarojgn810

PROFESSIONAL SUMMARY

Site Reliability Engineer with 9+ years designing, scaling, and operating mission-critical, high-availability distributed systems for Fortune 500 enterprises. Deep expertise across SRE, DevOps, AIOps, MLOps, and Agentic AI — building LLM-powered autonomous incident response, predictive anomaly detection, GenAI-driven runbook automation, and self-healing infrastructure on cloud-native Kubernetes platforms. Now expanding into a Forward Deployed Engineer track — training on the Microsoft / Azure AI stack (GitHub Copilot, Azure AI Engineer, Agentic AI Solutions Architect) and building and deploying end-to-end RAG-based GenAI applications. Proven track record of 99.9% uptime, 40% MTTR reduction, and large-scale toil elimination on payment-grade systems processing \$2B+ in daily transactions. AWS-certified and Executive PhD researcher in AI/ML at NIT Rourkela.

TECHNICAL SKILLS & CORE COMPETENCIES

SRE & Reliability	SLO / SLI / Error Budgets, Incident & On-Call Management, Chaos Engineering, Game Days, Root Cause Analysis (RCA), Blameless Postmortems, Toil Reduction, Capacity Planning, HA/DR
AIOps / MLOps / GenAI	AIOps Platform Engineering, LLM Runbook Automation, Agentic AI, LLM Orchestration, Predictive Anomaly Detection, ML Model Monitoring, Data & Model Drift, RAG, MLOps Pipelines
Cloud & Platforms	AWS (EC2, EKS, RDS, IAM, VPC, CloudWatch, Lambda, S3), Kubernetes, Docker, Helm, Service Mesh
IaC & CI/CD	Terraform, CloudFormation, Ansible, Jenkins, GitHub Actions, GitOps, ArgoCD, Release Automation
Observability	Prometheus, Grafana, Datadog, Splunk, New Relic, Elasticsearch / ELK, OpenTelemetry, Distributed Tracing, APM
Programming & Automation	Python, Bash / Shell, Java, Spring Boot, REST APIs, SQL, Autosys, Cron Automation
OS & Networking	Linux (RHEL / CentOS / Ubuntu), TCP/IP, DNS, Load Balancing, Firewalls, Networking Fundamentals
Methodologies	Agile / Scrum, ITIL, SRE Practices, DevSecOps, Change Management, SLA Compliance, Cross-functional Global Delivery

PROFESSIONAL EXPERIENCE

Site Reliability Engineer

Feb 2024 – Present

Wipro

Bangalore, India

- **AIOps Platform:** Architected AI-driven alert correlation across a Splunk and New Relic observability stack, reducing alert noise by 55% and dramatically improving on-call signal-to-noise ratio.
- **LLM Runbook Automation:** Deployed LLM-assisted triage automation for 30+ recurring incident types, reducing mean time to triage (MTTT) by 25% and eliminating manual first-response steps.
- **SLO Engineering:** Designed an SLO/SLI/error-budget framework for a payment gateway processing \$2B+ in daily transactions, sustaining a 99.9% uptime SLA across all production environments.
- **MTTR & Automation:** Built AI-powered automated remediation that reduced MTTR by 40%, saving 15+ engineering hours per month on critical payment incidents.
- **Predictive Scaling:** Deployed ML-based capacity models forecasting traffic spikes; proactive auto-scaling prevented 12+ potential SLA breaches during peak transaction windows.
- **Performance Engineering:** Resolved transaction bottlenecks through query optimization and ML-assisted load analysis, cutting payment failure rate from 2.3% to 0.8% (65% improvement).
- **Agentic AI (Active Research):** Architecting a multi-agent LLM system for autonomous incident response — AI agents that detect anomalies, execute runbooks, perform RCA, and trigger self-healing actions integrated with Kubernetes and Prometheus.

Principal Site Reliability Engineer

Jan 2023 – Dec 2023

Accion Labs India Pvt. Ltd.

Bangalore, India

- **Intelligent Incident Management:** Introduced AI-driven incident classification and routing across 150+ production applications, achieving 98% first-contact resolution and reducing escalations by 40%.
- **MLOps Observability:** Built MLOps pipelines monitoring model health, data drift, and prediction latency, reducing undetected model degradation by 70%.

- **Toil Elimination:** Automated 15+ SRE workflows using Python and AI-assisted scripting, cutting manual toil by 65% and reclaiming 25+ engineering hours per week.
- **Kubernetes SRE:** Operated Kubernetes clusters with AI-enhanced deployment gates, reducing deployment failures from 15% to 2% and improving release frequency by 50%.
- **Observability Platform:** Centralized logging with ElasticSearch, Grafana, and OpenTelemetry distributed tracing; intelligent alert escalation cut mean time to notify (MTTN) by 45%.

Career Break — Professional Development & Family

Dec 2019 – Dec 2022

Self-Directed Learning

Bangalore, India

- Took a planned career break for family responsibilities while proactively upskilling in emerging technology domains to prepare for a return to hands-on engineering.
- Completed self-directed study in Product Management and Artificial Intelligence / Machine Learning, building foundations later applied directly to AIOps and MLOps engineering work.
- Studied cloud architecture, Kubernetes, and modern SRE best practices, and earned AWS Cloud and AI Practitioner certifications during this period.

Specialist — Infrastructure & Operations

May 2017 – Nov 2019

Merck

Bangalore, India

- Managed 25 virtual collaboration platforms for 55+ enterprise clients across 12 countries, lifting availability from 94% to 99.2% and reducing support ticket volume by 30% through proactive monitoring and self-service automation.
- Spearheaded digital transformation initiatives (Regulatory Forecaster Tool, China RTR Digital Tool) that improved operational efficiency by 40%, earning an End-of-Year Performance Award.
- Led incident response and change management for globally distributed infrastructure, coordinating across multiple time zones to meet strict enterprise SLAs.

Software Engineer — L1/L2 Support

Jan 2014 – Apr 2017

PS Networks

Bhubaneswar, India

- Maintained 99.5% SLA compliance for enterprise batch processing and automated 25+ Autosys jobs, eliminating 20+ manual hours per week (80% toil reduction).
- Mentored 5+ junior engineers and improved team MTTR by 40% through structured SRE training and knowledge transfer.
- Performed root cause analysis and production troubleshooting across Linux systems, scripting, and job scheduling to sustain reliable operations.

PROJECTS

Enterprise RAG Chat Engine — Capstone Project

2025 – Present

Personal / Self-Directed

Bangalore, India

- Designed, built, and deployed an end-to-end Retrieval-Augmented Generation (RAG) chat engine for Banking and E-commerce use cases, enabling secure natural-language querying over domain-specific knowledge bases.
- Implemented the full pipeline — document ingestion, embeddings, vector search, LLM orchestration, and a conversational UI — and deployed to the cloud with CI/CD for reliable, repeatable releases.
- Applied Azure AI and GitHub Copilot tooling from the Forward Deployed Engineer training curriculum to accelerate development, testing, and delivery.

EDUCATION

Executive PhD — Artificial Intelligence & Machine Learning (Pursuing), National Institute of Technology (NIT) Rourkela, India

2026 – Present

B.Tech — Electronics Engineering, Krupajal Engineering College, Bhubaneswar, India

2007 – 2011

CERTIFICATIONS & TRAINING

AWS Certified AI Practitioner • AWS Certified Cloud Practitioner • ITIL Foundation • Agile / Scrum • Linux Administration Fundamentals

In Progress (Microsoft / Forward Deployed Engineer track): GH-300 GitHub Copilot • AI-103 Azure AI Engineer Associate • AB-100 Agentic AI Business Solutions Architect