

ALEX MERCER

Senior Site Reliability Engineer

✉ alex.mercer@nexa-sre.com | 📞 +1 (555) 489-0192 | 🌐 San Francisco, CA (Remote)
🐙 github.com/alex-mercere-sre | 🔗 linkedin.com/in/alex-mercere-sre

PROFESSIONAL SUMMARY

Performance-driven Senior Site Reliability Engineer with over 7 years of experience designing, building, and scaling high-availability cloud infrastructure. Proven track record of maintaining 99.99% uptime for critical microservice platforms. Expert in container orchestration (Kubernetes), Infrastructure as Code (Terraform), and implementing robust observability pipelines. Passionate about reducing MTTR, optimizing cloud expenditure, and cultivating a proactive post-mortem and blameless post-incident culture.

TECHNICAL CORE COMPETENCIES

Cloud & IaC: AWS, Nexa Cloud Platform (GCP), Terraform, Terragrunt, Ansible, Packer
Containers & Orchestration: Kubernetes (EKS/GKE), Docker, Helm, Istio Service Mesh, Nomad
Observability & Telemetry: Prometheus, Grafana, Datadog, ELK Stack, OpenTelemetry, Jaeger
Systems & Networking: Linux (Ubuntu/RHEL), Bash, Python, Go, TCP/IP, DNS, BGP, gRPC, Nginx
CI/CD & GitOps: GitLab CI, GitHub Actions, ArgoCD, Spinnaker, Artifactory

PROFESSIONAL EXPERIENCE

Senior Site Reliability Engineer

Nexa Cloud

Jan 2024 – Present

San Francisco, CA (Remote)

- Designed and executed a multi-region active-active failover architecture on AWS using Terraform and Route 53, reducing failover recovery time (RTO) from 45 minutes to under 90 seconds.
- Established service SLOs and error budgets for 15+ microservices, leading to a 40% reduction in alert fatigue and improving overall platform availability from 99.9% to 99.99%.
- Authored a Go-based Kubernetes controller to automate the dynamic scaling of node groups based on Prometheus custom queue-depth metrics, reducing average monthly infrastructure spend by \$10,000.

Site Reliability Engineer II

Vertex Systems

Mar 2021 – Dec 2023

Austin, TX (Hybrid)

- Successfully migrated legacy monolithic applications to Kubernetes (GKE), containerizing 24 core services and improving overall server resource utilization by 50%.
- Implemented GitOps deployment pipelines using GitLab CI and ArgoCD, reducing application deployment cycle times by 70% while securing continuous canary deployments.
- Developed and maintained centralized monitoring dashboards in Grafana with complex Prometheus queries, ingestion rates exceeding 5 million metrics/sec across 300+ nodes.

Systems & Infrastructure Engineer

Synapse Technologies

Jun 2019 – Feb 2021

Seattle, WA

- Provisioned bare-metal and cloud servers using Ansible, reducing developer environment onboarding times from 3 business days to under 20 minutes.
- Administered high-throughput PostgreSQL and Redis database clusters, establishing replica lag monitoring and automated daily backup verification workflows.

KEY PROJECTS

ChaosNimbus | Go, Kubernetes API, Prometheus

Open Source Project

- Built and open-sourced a chaos engineering daemon that schedules latency injection and pod disruptions to validate cluster resiliency under simulated high-stress conditions.

AutoRunbook | Python, Alertmanager Webhooks, AWS Lambda

Internal Infrastructure Tool

- Developed a self-healing lambda function triggered by Alertmanager webhooks, resolving 65% of persistent disk space and service lock alerts without manual operations.

EDUCATION

B.S. in Computer Science

Meridian University

GPA: 3.8/4.0

2015 – 2019

CERTIFICATIONS

- Certified Kubernetes Administrator (CKA)
- AWS Solutions Architect – Professional
- HashiCorp Certified: Terraform Associate