JULIAN MERCER

SITE RELIABILITY ENGINEER

San Francisco, CA | (415) 555-0148 | julian.mercer@email.com | linkedin.com/in/julianmercer

PROFESSIONAL SUMMARY

Site Reliability Engineer with 8+ years running high-traffic services with SLOs, observability, and disciplined incident response. Reduces downtime, automates toil, and partners with developers on capacity and resilience before problems reach customers.

CORE SKILLS

SLOs and error budgets | Incident response | Observability | Prometheus | Grafana | Kubernetes | Terraform | Go | Python | Linux | Capacity planning | Chaos testing | Postmortems | PagerDuty | Distributed systems

PROFESSIONAL EXPERIENCE

Staff Site Reliability Engineer Apr 2021 - Present
Datadog, New York, NY

- Defined SLOs and error budgets for twelve customer-facing services, giving product teams a shared basis for balancing releases against reliability work.
- Led incident command for major outages and cut mean time to recovery from 74 minutes to 38 minutes through clearer roles, runbooks, and paging rules.
- Reduced alert noise by about 55% by retiring low-value pages and moving to symptom-based alerts tied to SLO burn rate.
- Built an automated failover test in Kubernetes that ran monthly and found three hidden single points of failure before any customer saw an impact.
- Mentored six engineers entering the on-call rotation and ran quarterly game days that raised confidence in the incident process.
- Partnered with backend teams on load tests before launches, setting capacity targets that kept p99 latency under 400 ms during a traffic spike of 2x normal volume.

Site Reliability Engineer Sep 2018 - Mar 2021
Dropbox, San Francisco, CA

- Improved availability of the file sync API from 99.90% to 99.97% through capacity planning, load shedding, and safer rollout practices.
- Wrote Go and Python automation that removed about 20 hours a week of manual operational toil from the team.
- Tuned autoscaling and queue limits ahead of seasonal peaks, so traffic grew 30% without a customer-visible degradation.
- Authored blameless postmortems for 25 incidents and tracked action items to closure, with 90% finished within the agreed deadline.
- Rolled out a tracing pipeline across the API tier so engineers could follow a slow request end to end, shortening root-cause analysis from hours to under 30 minutes.

Systems Engineer Jul 2016 - Aug 2018
Yelp, San Francisco, CA

- Operated Linux fleets and load balancers supporting search traffic, responding to alerts and restoring service within agreed targets.
- Created Grafana dashboards and log queries that cut the time to isolate a failing host from about 30 minutes to ten.
- Moved manual deploy steps into Ansible playbooks, which removed configuration drift across 120 servers.

EDUCATION

B.S. in Computer Engineering, Rutgers University, New Brunswick, NJ

CERTIFICATIONS

Certified Kubernetes Administrator (CKA) | AWS Certified Solutions Architect - Associate