SRE Quick Reference

From Chaos to Confidence

The Four Golden Signals

`Latency`	Time to service requests (p50, p95, p99)
`Traffic`	Demand on your system (requests/sec)
`Errors`	Rate of failed requests (5xx, timeouts)
`Saturation`	How full your service is (CPU, memory)

SLO Formulas

`Availability`	(successful requests / total) x 100
`Error Budget`	100% - SLO target (e.g., 0.1% for 99.9%)
`Burn Rate`	Error consumption / time elapsed
`Time to Exhaust`	Remaining budget / burn rate

Performance & Response

DORA Metrics

`Deploy Freq`	Elite: on-demand, multiple/day
`Lead Time`	Elite: <1 hour commit to prod
`Change Fail`	Elite: <5% of changes fail
`MTTR`	Elite: <1 hour to restore service

Incident Severity

`SEV1`	Critical: >50% users, page immediately
`SEV2`	Major: Feature down, notify on-call
`SEV3`	Minor: Degraded, business hours
`SEV4`	Low: Cosmetic, queue for sprint

> **Remember:** Elite teams achieve BOTH high velocity AND high stability

Reliability & Practice

Cost of Nines

`99%`	3.65 days downtime/year ($)
`99.9%`	8.76 hours downtime/year ($$)
`99.95%`	4.38 hours downtime/year ($$$)
`99.99%`	52 minutes downtime/year ($$$$)

Chaos Engineering

`Hypothesis`	Define expected steady state
`Experiment`	Inject real-world failures
`Minimize`	Limit blast radius carefully
`Learn`	Improve resilience from findings