

Yuki Watanabe

Senior Site Reliability Engineer

Portland, OR · yuki.watanabe@gmail.com · +1 503-555-0142 · linkedin.com/in/yukiwatanabe

PROFILE SUMMARY

- **Senior Site Reliability Engineer** with **7 years** on high-traffic SaaS across **CI/CD, identity, and multi-tenant control planes**, specializing in **SLO design, incident command, and reliability architecture**.
- Hands-on across primary language (**Go**), scripting (**Python**), metrics stack (**Prometheus**), tracing (**OpenTelemetry**), orchestration (**Kubernetes**), infrastructure as code (**Terraform**), and chaos tooling (**Gremlin**), with strong fundamentals in distributed systems.
- Deep expertise in **error-budget policy, alert quality, graceful degradation, and capacity planning**, using **blameless postmortems** and **game days** to ship changes that survive a bad Tuesday.
- Acts as **incident commander** on major events, pairing with product engineering, support, and security across a fully remote org, running severity calls, comms, and the review that follows.
- Sets the reliability bar for the team: owns the on-call handbook, mentors engineers through their first rotations, and keeps the postmortem backlog from turning into a graveyard.

TOOLS & SKILLS

Languages:	Go, Python, Bash, SQL, PromQL
Observability:	Prometheus, Grafana, OpenTelemetry, Loki, Tempo, structured logging, SLI/SLO instrumentation
Reliability:	error-budget policy, incident command, blameless postmortems, graceful degradation, capacity planning
Resilience Testing:	Gremlin, Chaos Mesh, game days, load and soak testing, failure injection
Platform:	Kubernetes, Terraform, Helm, Argo CD, service mesh, progressive delivery
Cloud & Data:	GCP and AWS, PostgreSQL, Redis, Kafka, object storage, multi-region failover
Ways of Working:	on-call handbook ownership, mentoring, reliability reviews, cross-team postmortem follow-through

EDUCATION

Portland State University	B.S. in Computer Science	Portland, OR · Sep 2015 - Jun 2019
---------------------------	--------------------------	------------------------------------

WORK EXPERIENCE

GitLab	Senior Site Reliability Engineer	Remote · Mar 2022 - Present
• Own reliability for the CI runner control plane, a multi-tenant system serving thousands of concurrent jobs, from SLO definition through incident response and the architecture changes that follow. • Rewrote the SLO set around user-visible journeys rather than host metrics and wired the error budget into release gating, taking availability from 99.5% to 99.95% across four quarters. • Rebuilt alerting on Prometheus and OpenTelemetry with symptom-based rules and burn-rate windows, cutting pages by 62% while shortening median time to detect. • Introduced quarterly game days with Gremlin, injecting dependency and zone failures against real traffic shadows, which surfaced 9 latent single points of failure before customers found them. • Act as incident commander on Sev1 and Sev2 events and run the review afterwards, driving follow-up actions to completion and holding repeat-incident rate at under 8%.		
Auth0 (Okta)	Site Reliability Engineer	Portland, OR · Aug 2019 - Feb 2022
• Ran production for identity services in Kubernetes across three regions, covering capacity, rollout safety, and a shared on-call rotation with the owning product teams. • Led the migration to Terraform-managed infrastructure with peer-reviewed plans and drift detection, removing the hand-edited console changes behind two of the previous year's outages. • Built the multi-region failover runbook and proved it in a live exercise, bringing tested recovery time down to under 12 minutes.		