

Esteban Morales

Junior Site Reliability Engineer

New York, NY · esteban.morales@gmail.com · +1 646-555-0188 · linkedin.com/in/estebanmorales

PROFILE SUMMARY

- **Junior Site Reliability Engineer** with **2 years** on production systems across **payments and identity services**, specializing in **observability, on-call response, and toil automation**, after two years in systems engineering.
- Hands-on across primary language (**Python**), metrics stack (**Prometheus**), dashboards (**Grafana**), orchestration (**Kubernetes**), infrastructure as code (**Terraform**), and incident tooling (**PagerDuty**), with working knowledge of Go and distributed tracing.
- Growing expertise in **SLI instrumentation, alert tuning, and runbook authoring**, using **blameless postmortems** and **error-budget reviews** to turn repeat pages into **fixes that hold**.
- Carries a **primary on-call rotation** alongside three senior engineers, working with product teams, platform engineering, and support during incidents, joining daily stand-ups, incident reviews, and release checks.
- Systems background built the useful half of the job: **can read a stack end to end under pressure**, then write the runbook so the next person on call does not have to.

TOOLS & SKILLS

Languages:	Python, Go (basics), Bash, SQL, YAML
Observability:	Prometheus, Grafana, Loki, OpenTelemetry (basics), SLI instrumentation, alert tuning
Reliability:	SLOs and error budgets, on-call rotation, runbook authoring, blameless postmortems
Platform:	Kubernetes, Docker, Terraform, Helm, Linux administration
Automation:	toil reduction scripting, cron and job hygiene, GitHub Actions, config management
Incident Tooling:	PagerDuty, Statuspage, incident channels, severity triage
Cloud:	AWS (EC2, EKS, RDS, CloudWatch), basic networking and IAM

EDUCATION

CUNY Hunter College	B.S. in Computer Science	New York, NY · Sep 2018 - May 2022
Linux Foundation	Certified Kubernetes Administrator (CKA)	Remote · Nov 2024

WORK EXPERIENCE

Coinbase	Junior Site Reliability Engineer	New York, NY · Feb 2024 - Present
• Hold a primary on-call rotation for two payment-facing services alongside three senior engineers, triaging pages, running severity calls, and writing the postmortem for every incident carried. • Instrumented SLIs in Prometheus for latency, availability, and error rate across 6 services, publishing Grafana dashboards the product teams now open in their own reviews. • Cut alert volume by tuning thresholds against real error budgets and deleting rules nobody acted on, taking the rotation from 34 pages a week to 9 without losing a real incident. • Automated the certificate rotation and log-retention chores that ate the on-call day, scripting them in Python against the AWS APIs, recovering roughly 6 engineer-hours a week. • Wrote the first runbooks for 11 services that had none, pairing with the owning teams on failure modes and recovery steps, and folded them into the incident channel template.		
Plaid	Junior Systems Engineer	New York, NY · Jul 2022 - Jan 2024
• Ran Linux and Kubernetes environments for internal platform teams, handling capacity, patching, and access, with an on-call share for the cluster itself. • Migrated environment provisioning from hand-built hosts to Terraform modules with peer-reviewed plans, taking a new environment from two days to under an hour. • Sat in on incident calls to take notes and chase actions, which is where the reliability side of the job came from and why the SRE pivot stuck.		