Ethan Brown Junior Performance Engineer

San Jose, CA · ethan.brown@gmail.com · +1 408-555-0124 · linkedin.com/in/ethanbrown

PROFILE SUMMARY

- **Junior Performance Engineer** with **2 years** load testing and profiling **video meeting services**, after a year as a **back-end Python engineer** on a marketplace listings API.
- Writes load tests in **k6** and **Locust**, profiles Python and Go services with **py-spy** and **pprof**, and reads latency in **Datadog** and **Grafana** on **Kubernetes** and AWS.
- Learning **workload modeling** and **capacity math** from senior engineers, and reports **p95 and p99** on every test instead of averages.
- Works with **platform, SRE and media teams** in weekly test reviews, release readiness checks and the **on-call shadow rotation**.
- Keeps the team's **k6 script library** tidy, wrote the **how to run a soak test** page in Confluence, and pairs with new hires on their first test.

TECHNICAL SKILLS

Languages & Scripting:	Python, Go, JavaScript, SQL, Bash
Load Testing Tools:	k6, Locust, JMeter (basics)
Profilers & Runtime Tools:	py-spy, pprof, flame graphs, Chrome DevTools
APM & Observability:	Datadog, Grafana, Prometheus, OpenTelemetry (basics)
Databases & Caching:	PostgreSQL, Redis, EXPLAIN ANALYZE
Cloud, Containers & CI/CD:	AWS (EC2, EKS), Kubernetes, Docker, GitHub Actions
Front-End Performance Tools:	Lighthouse, WebPageTest, Core Web Vitals

EDUCATION

San Jose State University — **B.S. in Computer Science** — San Jose, CA · Aug 2019 to May 2023

WORK EXPERIENCE

Zoom — **Performance Engineer I** — San Jose, CA · Aug 2024 to Present

- Wrote **k6** load tests for the meeting join API from a week of production request logs, holding a **1,200 requests per second** baseline that the platform team now checks before each release.
- Profiled the recording upload workers with **py-spy** and found a JSON serialization hot spot, cutting CPU use at steady load from 71% to **48%**.
- Tuned 6 slow **PostgreSQL** queries behind the meeting history page under a senior's review, adding 2 indexes and bringing p95 query time from 620 ms to **140 ms**.
- Added a nightly **k6** smoke test to **GitHub Actions** that fails when join latency rises over budget, catching **5 regressions** before staging in its first 6 months.
- Built p50, p95 and p99 latency panels in **Datadog** for 4 service teams and reported each one against its p99 budget weekly, keeping join latency in budget for **11 of 12 months**.

eBay — **Software Engineer** — San Jose, CA · Jun 2023 to Jul 2024

- Ran **Lighthouse** audits on the listing page and moved image resizing to the CDN, cutting 75th percentile LCP from 3.9 s to **2.4 s**.
- Wrote a **Locust** test for the search API before a holiday sale, which found a connection limit that capped it at **900 requests per second**.
- Planned the listings cache sizing for the sale with a senior engineer from a year of traffic data, so it kept up with **2.5x** normal load.