Solenn Marek Lead Performance Engineer

Burbank, CA · solenn.marek@gmail.com · +1 818-555-0149 · linkedin.com/in/solennmarek

PROFILE SUMMARY

- **Lead Performance Engineer** with **11 years** on **video streaming platforms**, leading **playback latency and live event readiness** for services with tens of millions of subscribers.
- Hands-on with load testing (**Gatling**, **k6**, traffic replay), APM (**Dynatrace**, Datadog), profiling (**async-profiler**, pprof), data (**Apache Cassandra**, Amazon DynamoDB) and **AWS**.
- Deep expertise in **multi-CDN capacity models**, **per-event headroom planning**, **latency SLOs across teams** and **release gates**, using **Little's Law** to size what each event needs.
- Works with **playback, content delivery, SRE and product teams** on event readiness reviews, design reviews and the incident bridge.
- Leads **8 performance engineers** across 4 squads, runs the **platform performance guild**, and wrote the org's **performance engineering career framework**.

TECHNICAL SKILLS

Languages & Scripting:	Java, Go, Python, TypeScript, SQL
Load Testing Tools:	Gatling, k6, JMeter, Go traffic replay tools
Profilers & Runtime Tools:	async-profiler, Java Flight Recorder, pprof, eBPF
APM & Observability:	Dynatrace, Datadog, OpenTelemetry, Grafana
Databases & Caching:	Apache Cassandra, Amazon DynamoDB, PostgreSQL, Redis
Cloud, Containers & CI/CD:	AWS (EC2, EKS, CloudFront), Akamai, Fastly, Kubernetes, Jenkins
Front-End Performance Tools:	Lighthouse, WebPageTest, real-user monitoring, video start time metrics

EDUCATION

UCLA — **B.S. in Computer Science** — Los Angeles, CA · Sep 2011 to Jun 2015

WORK EXPERIENCE

Disney Streaming — **Lead Performance Engineer** — Burbank, CA · Mar 2021 to Present

- Lead 8 performance engineers across 4 playback squads, owning the latency SLOs and the readiness review for every live event, with **0 SEV1 incidents** across 38 events.
- Modeled CDN and origin capacity for 6 major live sports and premiere events with **Little's Law** and past viewer curves, sizing each one for **5.2x** normal traffic.
- Replaced legacy JMeter rigs with **Gatling** and a Go traffic replay tool, retiring 120,000 lines of test code and cutting full test runs from 9 hours to **2.4 hours**.
- Ran the cross-squad playback start program with **Dynatrace** traces and video client metrics, cutting p99 time to first frame from 1.6 s to **480 ms**.
- Made latency checks a required gate in the **Jenkins** release pipelines of 5 product orgs, with budgets agreed in RFCs, blocking **47 slow releases** in 2025.

Hulu — **Senior Performance Engineer** — Santa Monica, CA · Jul 2015 to Feb 2021

- Profiled the recommendations API with **async-profiler** and fixed thread pool starvation, keeping p99 under 110 ms at **6,000 requests per second**.
- Tuned **Apache Cassandra** read paths for the watch history service, moving to a time-bucketed schema that cut p95 reads from 90 ms to **18 ms**.
- Cut home screen load time on smart TV apps from 4.8 s to **2.6 s** by trimming image sizes and deferring rows below the fold, measured with real-user timing.