Amara Ndiaye Senior Go Developer

Paris, France · amara.ndiaye@gmail.com · +33 6 55 01 92 47 · linkedin.com/in/amarandiaye

PROFILE SUMMARY

- **Senior Go Developer** with **8 years** on **high-traffic marketplace and cloud services**, now owning **buyer-seller messaging and search indexing** at a **French classifieds site**, after 3 years on **Kubernetes and CLI tooling** at a cloud provider.
- Builds in **Go** with **gRPC** and REST, keeps data in **PostgreSQL** and **Redis**, streams with **Apache Kafka**, runs on **Kubernetes** and watches it all in **Prometheus** and OpenTelemetry.
- Strong in **concurrency design** (bounded worker pools, backpressure, cancellation), **memory and GC tuning** and **test suites** with fuzzing and the race detector, and backs each performance change with **benchstat** numbers.
- Works with **product managers**, the **search and trust and safety teams**, **SRE** and **mobile engineers**, and takes the **messaging on-call rotation** one week in five.
- Leads **design reviews** for a team of 7, wrote the team's **concurrency review checklist**, and mentors **2 engineers** who moved to Go from PHP and Java.

TECHNICAL SKILLS

Languages:	Go, SQL, Bash, Python, PHP (legacy)
Go Libraries & Frameworks:	net/http, chi, Cobra, client-go, sqlc, pgx, segmentio/kafka-go
APIs & Protocols:	gRPC, Protocol Buffers, REST, OpenAPI, WebSockets
Databases & Messaging:	PostgreSQL, Redis, Apache Kafka, Elasticsearch
Cloud, Containers & CI/CD:	Kubernetes, Docker, Helm, Argo CD, Terraform, GitLab CI
Testing & Code Quality:	go test, testify, Testcontainers, native fuzzing, benchstat, golangci-lint
Observability & Profiling:	pprof, runtime/trace, Prometheus, OpenTelemetry, Grafana, Pyroscope

EDUCATION

Télécom Paris — Diplôme d'ingénieur (Master's level), Computer Science — Paris, France · Sep 2015 to Jun 2018

WORK EXPERIENCE

leboncoin — Senior Backend Engineer (Go) — Paris, France · Apr 2021 to Present

- Moved buyer-seller messaging off a legacy service onto a **Go** service with **gRPC** inside the platform and a REST edge for the apps, cutting p99 latency from 340 ms to **55 ms**.
- Redesigned ad moderation as an **errgroup** fan-out with bounded worker pools and per-check **context** deadlines, so one slow image check no longer stalls a batch, raising throughput to **4,200 ads a minute**.
- Built the **Apache Kafka** consumer group that feeds new ads to search indexing, with idempotent **PostgreSQL** writes and lag-based backpressure, cutting indexing delay from 6 minutes to **under 20 seconds**.
- Cut memory on the search gateway with **pprof** heap profiles, generated JSON decoders instead of reflection, sync.Pool buffers and a **GOMEMLIMIT** per pod, so it now runs on **40% fewer pods** at the same traffic.
- Set up **Testcontainers** suites for Kafka and PostgreSQL, native fuzzing on the ad parser and **-race** in every CI run across 6 services, catching **17 concurrency bugs** before production in a year.

Scaleway — Software Engineer — Paris, France · Jul 2018 to Mar 2021

- Wrote **14 commands** for the open source **Scaleway CLI** in Go (instances, object storage, Kubernetes clusters), with generated docs and release builds for Linux, macOS and Windows.
- Extended the Kubernetes **cloud controller manager** in Go with **client-go** so load balancer settings follow Service annotations, cutting support tickets about misconfigured load balancers by **60%**.
- Moved **9 repositories** from GOPATH and dep to **Go modules** in 2019, pinning versions and adding a weekly upgrade job, which cut CI build time from 14 minutes to **6**.