

Hi,

I'm Burak. Twelve years in this game, currently running a one-person DevOps shop out of Istanbul (BRF Tech). Most of that time I've spent elbow-deep in other people's production. Some of it I fixed. A fair chunk I rebuilt because nobody on the team wanted to admit it should have been retired five years ago.

Stuff from the last two years that's probably relevant to whatever you're hiring for:

At Monotech (iGaming, Zagreb-based, I worked remote from here) I took a monolith they had on GCP Serverless and moved it onto GKE. End result was roughly half the infra cost gone, while we kept serving the flagship customer's 100k–500k users plus three to five other tenants. Per-customer GKE clusters that autoscale 20–30 nodes. Cloudflare WAF custom rules and rate limiting in front. Replaced their Vault setup with Passbolt + External Secrets, which let us push credentials cluster-wide from a single source of truth. Built a self-service Filament dashboard so devs could tweak their own WAF rules instead of opening an ops ticket every other day.

Before Monotech, back at C-Level, I built PaketTaksi OYM. It's a serverless reconciliation job. Pulls roughly 50M fuel records out of Petrol Ofisi's API every night at 23:55 and reconciles them per courier in 1 to 3 minutes on AWS Lambda + Vapor. SLA was five. The thing it replaced was an Excel + macro nightmare that took someone an hour every morning. I designed it, wrote it, shipped it, and ran it for years afterwards.

Security stuff I actually do, not just list on a slide. Cloudflare WAF custom rules. OWASP-aware code review (the kind where you read the diff, not the kind where you run a tool). Image scanning with Trivy. Vault and gitleaks pre-commit hooks for secrets. Kubernetes RBAC, Network Policies, Pod Security. mTLS through cert-manager. OIDC and SAML integrations against Keycloak, Authelia, or Auth0 depending on what the customer already uses. Audit logging into SIEM. KVKK and GDPR work where the compliance team actually needed it.

On the side I'm cooking runr.sh — open-source serverless on Knative, with a native AI module (single-agent execution, multi-agent DAG pipelines, GPU-pooled model serving with an OpenAI-compatible inference API). Pre-release, sole engineer. Built specifically for fintech and government deployments where data residency and confidential computing matter. Plus a couple of upstream Linux kernel patches landed (ThinkBook ACPI, Lenovo fingerprint sensor driver), and a self-hosted production lab at home that runs about a dozen services on a Kubernetes 1.24 cluster.

Two CVs attached:

- **Burak_Sahin_CV.docx** — the compact one. Two pages. ATS-friendly, every keyword you might be filtering on.
- **Burak_Sahin_CV_Detailed.docx** — the long one. Five pages, with the actual stories behind each role. What was hard, what I built, what I learned. Read this if you want technical depth.

I'm looking at senior, staff, lead, principal IC, or CTO-track work. Remote (EU or US) is fine. Hybrid is fine. I'd also consider relocating for the right team. If anything in here matches what you're trying to hire for, ping me. Happy to jump on a call and walk through any of these systems, share architecture notes, or just talk about whatever problem your team is currently stuck on.

Thanks for reading this far.

— Burak

Founder, BRF Tech · brf.sh · burak@brf.sh · burak.sh