Anna Sbrodova

Location: Indonesia (WITA / UTC+8) **| Phone:** +62 813 3793 8028 **| E-Mail:** mail@asbrodova.dev

Personal Portfolio Website | GitHub (Open-Source Showcase) | Technical Medium Blog

Senior Software Engineer / Principal Backend Engineer (Go, GCP)

Backend Systems Architect | Distributed Systems Specialist | High-Load & Scalability Engineering | International Remote Delivery

Technically proficient and product-minded professional with 12+ years of software engineering experience across enterprise platforms, distributed backend systems, cloud-native infrastructure, and event-driven architectures within fintech, logistics, automation, and large-scale e-commerce environments. Over the past several years, specialised in designing and scaling high-throughput Go/GCP backend systems focused on distributed processing, reliability engineering, and cloud-native platform delivery. Strong end-to-end ownership capability, able to take complex backend products from stakeholder requirements and system design through architecture, implementation, infrastructure, observability, and production. Experienced in working independently on complex technical problems and translating product and business requirements into well-structured backend architectures.

KEY QUALIFICATIONS

- Go (Golang) Backend Development
- Google Cloud Platform (GCP)
- Distributed Systems Design
- Event-Driven Architecture
- Microservices Architecture
- Kafka Streaming & Messaging Systems
- Cloud Run / Serverless Architecture
- High-Load System Design (Scalability & Throughput)
- System Performance Optimisation
- CI/CD Pipeline Automation
- Architectural Thinking & Decision-Making
- End-to-End Ownership & Accountability

KEY CAREER ACCOMPLISHMENTS

- Architected and delivered a Go/GCP event-driven processing platform handling tens of thousands of Kafka events under peak load, ensuring stable throughput and fault-tolerant execution across distributed backend services.
- Designed large-scale distributed compute pipelines orchestrating hundreds of parallel Cloud Run jobs, each processing millions of BigQuery rows per execution for high-volume data transformation workloads.
- Built an enterprise cloud expenditure chargeback platform from scratch, translating complex multi-tenant cloud billing and SaaS cost models into scalable automated business logic and production-ready distributed workflows.
- Supported large-scale e-commerce infrastructure at ManoMano handling ~50M monthly visits and scaling from 3,000-5,000 concurrent users to 15,000-30,000+ users during Black Friday peak events.
- Solved critical distributed consistency challenges by implementing Transactional Outbox and Inbox patterns, eliminating dual-write vulnerabilities and ensuring resilient, idempotent event delivery across business-critical financial systems.
- Reduced BigQuery operational expenditure by ~18% through query optimisation, schema redesign, and execution-flow improvements, while also improving Cloud SQL infrastructure efficiency and deployment automation across cloud-native environments.
- **Aura Tracker GCP, an open-source Go-based Model Context Protocol (MCP) server:**
  - Built a commercial-grade foundation Go-based MCP server exposing 67+ infrastructure tools across 24 GCP domains, enabling secure AI-driven cloud orchestration. Designed a strict Hexagonal Architecture with rate-limited, async execution (token-bucket + errgroup) and a pluggable module system for scalable tool expansion. Implemented enterprise-grade safety and observability layers including dry-run/safe-apply controls and PII anonymisation for secure LLM-integrated infrastructure access.

PROFESSIONAL EXPERIENCE

Hands-on Technical Lead ■ EPAM Systems Mar 2025 – Present

Design and deliver distributed backend systems on Go and GCP from initial stakeholder requirements through system design, implementation, deployment, and production hardening. Lead engineering across large-scale data pipelines, asynchronous processing systems, and multi-service distributed environments spanning Kafka, Pub/Sub, BigQuery, Cloud Run, and GKE.

Distributed Systems (Go / GCP) | Event-Driven Architecture | Kafka | Cloud-Native Platforms | High-Load System Design | Observability & Reliability Engineering | Platform Engineering | CI/CD Automation | Distributed Data Processing

SELECTED CONTRIBUTIONS

- Owned end-to-end architecture across cloud-native systems spanning Kafka, Pub/Sub, BigQuery, Cloud Run, and GKE, defining service boundaries and ensuring scalable system design across multi-service environments.
- Delivered and operated a 15+ microservice GKE ecosystem structured around event-driven patterns, ensuring loose coupling, independent service scalability, and production resilience.
- Optimised Cloud SQL performance through workload profiling and right-sizing of PostgreSQL instances.
- Implemented full-stack observability (Grafana, Loki, Prometheus, OpenTelemetry), improving system visibility and reducing time to diagnose production issues across distributed workflows.
- Built centralised CI/CD deployment orchestration tooling using automated YAML manifests capable of calculating dependency-aware microservice deployment sequencing without manual coordination.
- Introduced runtime dependency graph resolution for deployments, ensuring safe releases across complex distributed service topologies.
- Established staging-to-production parity monitoring, enabling early detection of regressions before production impact.
- Designed proactive alerting pipelines integrated with collaboration tools that improved real-time detection of system anomalies.
- Acted as architecture owner across multiple initiatives and translated product requirements into structured engineering designs with clear execution pathways for distributed systems.
- Mentored engineers through technical reviews focused on distributed systems design, event-driven architecture, and Hexagonal principles.
- Collaborated directly with neighbouring engineering teams to define clean API contracts and eliminate integration friction across distributed services and asynchronous workflows.

Senior Backend Engineer / Technical Lead ■ EPAM Systems

May 2023 – Mar 2025

Designed and delivered distributed backend systems on Go and GCP, supporting enterprise-scale workloads across cloud-native, event-driven architectures. Worked across multiple high-throughput services involving Kafka-based event processing, asynchronous workflows, and large-scale data pipelines. Focused on system scalability, reliability, and production-grade execution.

Distributed Systems (Go / GCP) | Event-Driven Architecture | Kafka | Cloud-Native Backend Systems | High-Load System Design | Observability & Reliability Engineering | Microservices Architecture | Scalable Data Processing

SELECTED CONTRIBUTIONS

- Built and maintained Go-based distributed microservices on GCP supporting enterprise-scale workloads across multiple high-throughput backend systems, ensuring stable production delivery across services handling variable traffic and data volume spikes.
- Designed event-driven architectures using Kafka, Redis, and GCP services (Pub/Sub, Cloud Run), enabling scalable asynchronous processing, reliable message handling, and loosely coupled service communication across distributed backend systems.
- Translated complex business and integration requirements into scalable backend system designs with clear service boundaries, well-defined contracts, and execution-ready architectural patterns aligned with production constraints.
- Introduced structured observability practices (metrics, logs, tracing) across services, improving system visibility, reducing blind spots in distributed workflows, and significantly lowering time required to isolate and diagnose production issues.
- Established baseline monitoring and alerting patterns aligned with service-level objectives (SLOs), improving early detection of performance degradation, bottlenecks, and failure conditions across distributed systems.
- Delivered hands-on engineering execution across multiple parallel initiatives while contributing to architectural decisions impacting scalability, resilience, and long-term maintainability of backend systems.

Senior Software Engineer ■ ManoMano

Feb 2022 – Apr 2023

Built, scaled, and optimised distributed systems in a fast-paced e-commerce environment, delivering production-grade backend services across Go and Java (Spring WebFlux) ecosystems. Collaborated with engineering and platform teams to support reliable system delivery across event-driven architectures, cloud-native infrastructure, and large-scale data processing pipelines

High-Scale Distributed Systems | Go & Java Backend Engineering | Spring WebFlux | Event-Driven Architecture | Performance Optimisation | Cloud-Native Systems (GCP) | High-Concurrency Systems | Backend Reliability Engineering | System Scalability & Load Handling

SELECTED CONTRIBUTIONS

- Supported production e-commerce systems handling 3,000-5,000 concurrent users (~50M monthly visits) across distributed Go-based backend services, ensuring stable service delivery under sustained traffic conditions.
- Scaled backend infrastructure during Black Friday peak events to 15,000-30,000+ concurrent users, maintaining system stability during extreme load spikes and unpredictable demand surges.
- Improved performance of critical Go services by removing bottlenecks in data access layers and stabilising high-concurrency transactional workflows under production load.

- Optimised backend execution paths through query tuning, caching improvements, and targeted profiling, resulting in consistently lower latency across high-throughput systems.
- Worked across event-driven systems using Kafka, Redis, MongoDB, and GCP services, supporting asynchronous processing and real-time data movement across services.
- Contributed to design and refinement of loosely coupled microservices aligned with event-driven architecture principles.
- Supported continuous improvement of backend reliability through production issue analysis, performance tuning, and incremental system hardening in live environments.

ADDITIONAL EXPERIENCE

Senior Software Engineer · IBA (Client: Hapag-Lloyd) Dec 2019 – Dec 2021

- Built Java/Spring Boot-based enterprise logistics and integration systems with asynchronous processing using IBM MQ in distributed backend environments.

Senior RPA Developer · Standard Bank (South Africa) Feb 2018 – Dec 2019

- Developed Java and Groovy-based automation solutions delivered from design to production within enterprise banking workflows.

RPA Developer · WF-TESSI Mar 2017 – Feb 2018

- Built Java and Groovy-based automation workflows integrating OCR and intelligent document processing into legacy enterprise systems.

Junior to Mid-Level Full-Stack Developer · Standard Bank & IBM 2014 – 2017

- Worked on Java and Spring-based enterprise banking systems supporting backend development and legacy platform maintenance.

EDUCATION

Master's Degree in Mechanics and Mathematics
Belarusian State University, 2014

CORE TECHNICAL STACK

Primary Languages & Cloud:	- Go (Golang), Google Cloud Platform (GCP), SQL
Cloud-Native Platforms:	- GKE (Google Kubernetes Engine), Cloud Run, Pub/Sub, BigQuery, Cloud SQL
Distributed Systems & Architecture:	- Event-Driven Architecture, Hexagonal Architecture (Ports & Adapters), Transactional Outbox Pattern, Transactional Inbox Pattern, Dual-Write Problem Resolution, High-Load System Design
Data & Messaging Systems:	- Kafka, Redis, PostgreSQL, MongoDB, IBM MQ
Infrastructure & DevOps:	- Terraform, Kubernetes, Docker, CI/CD automation, Eventarc
Observability & Reliability Engineering:	- Prometheus, Grafana, Loki, OpenTelemetry, GCP Monitoring, Structured Logging, Proactive Alerting Systems
Platform Engineering & Internal Tooling:	- CI/CD pipeline automation, dependency-aware deployment orchestration, microservice release automation, GitOps-style workflows
AI / MCP / Advanced Systems:	- Model Context Protocol (MCP), cloud orchestration tooling, AI-safe infrastructure APIs, secure LLM tool exposure patterns

--References available upon request--