

ANDREY VASILIEV

Senior Python Backend Engineer | Distributed Systems & High-Load Architecture

Saint Petersburg, Russia | andrew.vasilev03@gmail.com | a.vasilev@strkt.ru | +7 (923) 757-24-77 | t.me/neffus | github.com/N3ffus | leetcode.com/u/N3ffus (1,000+ Solved)

PROFESSIONAL SUMMARY

Senior Python Backend Engineer with **4+ years of commercial experience** specializing in distributed systems, asynchronous microservices, and database performance optimization. Engineered scalable platforms processing **40M+ records**, handled **20K+ daily events**, and reduced analytical query latencies from **60s down to <1.5s** with ClickHouse and PostgreSQL. Proficient in designing production Retrieval-Augmented Generation (RAG) platforms using pgvector, orchestrating event-driven architectures (FastAPI, Celery, Kafka/RabbitMQ), and enforcing automated CI/CD pipelines with **80%+ test coverage**.

WORK EXPERIENCE

Senior Python Backend Engineer | [Strikt](#) | 09/2024 – Present | Remote, Saint Petersburg

Real-Time Telegram Analytics & Ingestion Engine (PostgreSQL, ClickHouse, Redis, Celery)

- **Spearheaded distributed data ingestion streams** handling **20K+ daily messages** and **40M+ cumulative records** with strict transactional consistency.
- **Accelerated analytical query execution by 95% (from 30–60s to <1.5s)** through migrating OLAP reporting workloads from PostgreSQL to ClickHouse materialized tables.
- **Trimmed API response times by 90% (from ~8s to <250ms)** by resolving N+1 ORM query bottlenecks, restructuring composite indexes, and deploying Redis caching.

Enterprise AI Knowledge & Semantic Search Platform (FastAPI, pgvector, LLM Orchestration, Docker)

- **Pioneered an end-to-end RAG architecture** across **50K+ knowledge documents**, enabling sub-second (**<400ms**) vector retrieval with pgvector and semantic embedding pipelines.
- **Orchestrated multi-agent AI workflows** leveraging DeepSeek and Qwen models, mitigating retrieval hallucinations and securing **99.9% production availability**.

Enterprise Conference Management & Payment System (Django, DRF, Yookassa, PostgreSQL, Celery)

- **Constructed zero-downtime backend microservices** for an enterprise events platform (Vedomosti) supporting **10K+ active attendees** across **25+ concurrent conferences**.
- **Delivered 30+ REST endpoints and domain schemas** with strict OpenAPI contracts, automated validation schemas, and idempotent Yookassa transaction handling.
- **Elevated test automation coverage to 80%+** using Pytest suites and CI/CD quality gates, curbing regression incidents by 40%.

Engineering Leadership & Technical Governance

- **Cultivated team excellence** by mentoring **5+ backend developers** through code reviews and system design sessions, while steering **20+ technical interviews**.
- **Standardized Clean Architecture principles**, ADRs, and shared CI/CD blueprints, slashing new microservice onboarding cycles by 50%.

Python Backend Engineer | [Freelance / Contract](#) | 09/2022 – 08/2024 | Remote

- **Built high-throughput job aggregation engines** scraping and processing **20K+ listings daily** across multi-source job boards with 99% deduplication precision.
- **Engineered resilient distributed ETL pipelines** using FastAPI, Celery, Redis, and RabbitMQ, reducing data anomalies by 30% through automated backpressure controls and robust error handling.
- **Automated user notifications and interactive bots** via aiogram, serving 2,000+ daily active subscribers.

TECHNICAL SKILLS

Languages & Core: Python, SQL, Data Structures & Algorithms (1,000+ LeetCode problems solved, Codeforces)

Backend Frameworks: FastAPI, Django, Django REST Framework, RESTful APIs, OpenAPI, Clean Architecture, SOLID, Microservices, AsyncIO

Databases & Storage: PostgreSQL, ClickHouse (OLAP), Redis (Caching & Locking), pgvector, MongoDB, Query Optimization, Sharding & Replication

Messaging & Queues: Celery, RabbitMQ, Kafka, Event-Driven Architecture, Distributed Task Queues, Background Job Processing

AI & LLM Technologies: Retrieval-Augmented Generation (RAG), Vector Embeddings, Semantic Search, AI Agents, DeepSeek, Qwen LLM Integration

DevOps & Infrastructure: Docker, Docker Compose, Kubernetes, Linux, Git, GitLab CI/CD, Nginx, Prometheus, Grafana, Sentry, Yandex Cloud

LANGUAGES

English: B2 (Upper-Intermediate) – Professional Working Proficiency | **Russian:** Native