

Ilya Tomin

@ertiops • straxisrul@gmail.com

Baku, Azerbaijan · Remote / Relocation

Backend Engineer with 5 years of experience developing Python backend services, API platforms, CRM/ERP systems and distributed applications. Worked with PostgreSQL, event-driven systems, high-load APIs, integrations and production backend engineering.

Skills: Python, Django, FastAPI, aiohttp, SQLAlchemy, Pydantic, REST API, WebSocket, PostgreSQL, Redis, Kafka, RabbitMQ, Docker, Linux, Nginx, S3, GitLab CI, GitHub Actions, React, TypeScript, Pytest, Keycloak, OAuth2/OIDC, OpenAPI, Clean Architecture, DDD, Unit of Work, Outbox Pattern, DLQ, PgBouncer

Languages: English (B2), Russian (Native)

Experience

NDA - HRTech

Python Backend Engineer

03.2024 – 08.2026 (2 years 6 months)

- Designed service boundaries and RabbitMQ contracts for HRTech backend decomposition, coordinating and reviewing two developers' work. Following combined architecture and runtime changes, peak successful HTTP RPS grew from ~1,500 to ~8,000.
- Built and shipped a staff activity audit log for an HR platform using Kafka and a transactional outbox. Implemented asynchronous event storage in a separate PostgreSQL database, deduplication by unique UUIDv7 event IDs, a diagnostic dead-letter topic, and role-based access controls.
- Replaced a third-party Keycloak integration with an aiohttp client that reduced unnecessary dependencies and eased integration testing. Implemented user and service authorization and user-data updates.
- Set up parallel pytest-xdist execution for ~2,000 tests, cutting CI runtime from ~20 to ~1.5 minutes (~13×) with separate test tables per worker.
- Extracted ORM-independent Python Use Cases from Django Views/ViewSets (Clean Architecture), cutting module dependencies from ~321 to ~201 (about 37%); consolidated ~13 repeated user retrieval/permission checks and enabled database-free unit tests.
- Implemented role-based access with Keycloak OAuth2/OIDC and composite roles, plus time-limited S3 presigned URLs for file access.
- Investigated a PostgreSQL candidate-list query after bulk UPSERT with EXPLAIN ANALYZE. Repeated production measurements showed SQL Execution Time falling from ~17 s to ~100 ms and a switch from Seq Scan to an existing Index Scan after autovacuum thresholds were adjusted.

Text Analytics - LangTech

Backend Developer & ML Engineer

09.2021 – 02.2024 (2 years 6 months)

- Developed an NLP term extractor for educational-text complexity analysis (spaCy NER, reference dictionaries): ~5 s automated extraction versus ~20 min manual term search and checking per comparable text; linguists validated the automated results separately.
- Developed ruBERT-based genre classification models, achieving an F1 score of 0.83 on the validation dataset.
- Developed exam variant management and access rules for a Russian-language examination platform for international students.
- Developed a Chromium-based Windows application in Visual Studio, restricted to the exam-platform website; helped avoid a ~RUB 500,000 purchase of a third-party proctoring system.
- Helped introduce GitFlow and parallel development in a three-developer team: comparable work took ~2.5 instead of ~5 working days (~2× faster).

Education

Kazan Federal University

Master's Degree in Computational Linguistics