

Ivan Tsyganov

AI Platform Architect | Principal Backend Engineer

Neuchâtel, Switzerland · (+41) 78 250 00 05 · tsyganov.ivan@gmail.com

tsyganov-ivan.com · linkedin.com/in/tsyganovivan · github.com/tsyganov-ivan

Born: 9 July 1988 · Nationality: Russian · Permit: B

Languages: English (C1, daily working language) · French (B1) · Russian (native)

PROFESSIONAL SUMMARY

Software Architect with 15 years of experience spanning cybersecurity and distributed systems, now applied to building secure, production-grade AI platforms. Currently leading the design and delivery of a multi-tenant LLM platform — including agentic frameworks, MCP integration, and vector storage — at Constructor Group. Proven track record of building products from scratch, growing engineering teams, and establishing architectural governance.

EXPERIENCE

Constructor Group

Dec 2021 – present

constructor.tech · Schaffhausen, Switzerland · Fully Remote

EdTech and AI company building intelligent learning and research platforms for universities and enterprises worldwide.

Architect, Principal Backend Engineer · Oct 2024 – present

- Lead a distributed team of 10 software and ML engineers across 4 countries to design and deliver the Core Model Engine subsystem
- Establish architectural design review process and engineering-wide standards (API design, distributed tracing, error handling, CI/CD, security) adopted across all product teams
- Apply security-first architecture to multi-tenant LLM platform: dedicated prompt injection classifier, sandboxed tool execution with restricted permissions, and per-tenant permission scoping ensuring agents can only access authorized tools and data regardless of LLM behavior
- Platform powers university and enterprise tenants across 5 countries; 10 internal ML engineering teams build customer-facing products on top of these components (99.5% uptime SLA)
- Architected LLM-connectors — a unified multi-tenant interface supporting 10+ LLM providers, handling ~30K requests/day with streaming normalization across providers, automatic retries with provider-level error handling, per-tenant usage tracking with rate limits and quotas; distributed tracing of agent sessions via Langfuse
- Built Agentic Engine — an API-driven framework for building ReAct agents with tool execution via MCP, serving ~1K agent sessions/day
- Delivered MCP Engine — a service for configuring MCP instances, exposing built-in tools, wrapping existing APIs as tools, or acting as a proxy to external MCP servers
- Built Document Storage — a multi-tenant vector data service with pluggable embedding models via LLM-connectors and chunking strategies, indexing ~3M chunks with p95 search latency under 300ms
- Delivered Memory Service — a service enabling agents to store, manage, and search their own memory
- Designed and built a dual sync/async Model Engine SDK with OAuth2 auth, lazy-load object hierarchy, SSE streaming, LangChain integration, and MCP client factory — plus 7 more internal Python packages adopted by all 10 ML engineering teams across all backend services

Lead Developer / Senior Python Developer · Dec 2021 – Oct 2024

- Joined as a Senior Python Developer on a 2-person backend team; grew the team to 5 engineers and took on Tech Lead responsibilities overseeing frontend, QA, and DevOps teams
- Designed and implemented a multi-cluster Kubernetes platform for researchers (3 compute clusters in production, hot-pluggable architecture), enabling launch and customization of Docker containers with integrated VS Code development environments and hardware resource selection
- Built a Kubernetes operator (kopf + kubernetes-asyncio) managing 6 Custom Resource types with auto-generated CRD schemas from Pydantic models, asyncio semaphore-based concurrency control, and distributed Sentry trace propagation via K8s annotations
- Built an Nginx/Lua API gateway — 16 Lua modules implementing JWT validation, Redis-backed grant caching, dynamic multi-cluster routing, and WebSocket upgrade support
- Designed and implemented a Workflow component — a long-running pipeline engine for executing computations across different software/hardware configurations, including SSH-accessible machines and HPC clusters

Positive Technologies

ptsecurity.com · Moscow, Russia

Jan 2014 – Jan 2022

Enterprise cybersecurity company with 1,000+ employees, specializing in vulnerability assessment, incident response, and security analytics.

Tech Lead / Senior Python Developer

- Joined as Senior Python Developer; grew into Tech Lead managing a 10-person development team
- Rewrote the integration testing tool, reducing execution time from 2+ hours to ~15 minutes
- Built The Standoff from concept to production as sole backend engineer — a real-time cybersecurity competition platform featuring a scoring engine with dynamic task pricing, SLA calculation via interval-merging algorithm, declarative rules engine, and a dual-channel event system (WebSocket + Kafka, ~30 event types)
- Bootstrapped the next-generation Standoff365 platform — modular architecture with strict layering, async pub/sub event system, and append-only event sourcing for scoring
- Platform used at international events: Positive Hack Days (Moscow), Hack In The Box 2019 (Abu Dhabi), and online at standoff365.com

TimeWeb

timeweb.com · Saint Petersburg, Russia · One of the largest hosting companies in Russia (750,000+ clients)

Jun 2011 – Nov 2013

Python Developer

- Migrated billing stored procedures to Python (enabling test coverage), built IMAP integration workers and financial reporting tools, established code review process across the engineering team

Ukhta State Technical University

Ukhta, Russia · Part-time

Sep 2007 – May 2011

Programmer, Linux System Administrator

- Developed the university website (Drupal) with ERP integration; launched the remote education infrastructure (Moodle)

EDUCATION

Ukhta State Technical University, Ukhta

Master's engineering degree in Computer Science · Sep 2005 – Jul 2010

TECHNICAL SKILLS

ARCHITECTURE

Distributed systems, Microservices, Event-driven, Multi-tenant, System design, Kubernetes CRDs

AI / LLM

OpenAI, Anthropic, Google, Mistral, Cohere, Meta Llama, MCP, Agentic frameworks, RAG, Qdrant, Langfuse

DATA & MESSAGING

PostgreSQL, SQLAlchemy, Redis, RabbitMQ, Kafka

QUALITY & SECURITY

pytest, OWASP, OAuth2, OIDC, JWT, GDPR

CORE

Python, FastAPI, AsyncIO, Django, DRF, Celery, Pydantic, REST API, WebSockets

INFRASTRUCTURE

On-premises Kubernetes (data residency), Helm, Docker, Docker Compose, nginx, Linux, CI/CD

OBSERVABILITY

Prometheus, Grafana, OpenTelemetry, Sentry

METHODOLOGY

Agile, Scrum

TECHNICAL LEADERSHIP & COMMUNITY

- Moscow Python Conf program committee member; mentored ~30 engineers from proposal through delivery of conference talks
- Speaker at PyCon Russia, Moscow Python Conf, QA Fest (2014–2019) on Python security, testing, and DSLs; regular presenter at Constructor Group internal meetups