

Sviat Nahirnyi

sviatoslav.nahirnyi@gmail.com | [linkedin.com/in/sviat-nahirnyi](https://www.linkedin.com/in/sviat-nahirnyi) | sviat.info | London, UK
Forward Deployed Engineer | Data & AI Platforms | Presales

Experience

Co-Founder & Lead, Data & AI Platforms

Mar 2023 - Aug 2026

GRAPHIT

London, UK

- Closed the company's largest enterprise deal - a top-20 European logistics company - by leading the pitch and end-to-end solution design
- Embedded on-site at the client's warehouses to map operational workflows, then delivered analytical data platform solutions that improved operational performance by 6%
- Worked forward-deployed at Dow Jones's Barcelona office, working on vector database scalability for production retrieval workloads
- Grew the Data & AI practice from zero to 17 engineers across 9 client engagements - real-time platforms, Lakehouse migrations, LLM evaluation systems
- Cut flagged chatbot responses by 30% by architecting an LLM-as-a-Judge evaluation platform

Senior Data Engineer

Jul 2023 - Jan 2025

SIGMA SOFTWARE GROUP

Munster, Germany

- Delivered the source-of-truth Lakehouse for an international logistics company - 39 operational sources consolidated into one platform serving 5 analytics and ML teams - by embedding on-site at the client's factory to design ETL flows with operations staff
- Designed analytical dashboards for factory and logistics stakeholders on-site, mapping shop-floor workflows directly into data models and reporting
- Reduced end-to-end pipeline runtime by 37% by building incremental PySpark pipelines on Databricks

Presales Software Engineer

May 2023 - Aug 2023

GREENM

Hybrid, US

- Won a data-platform tender for an aerospace company via a one-month PoC delivered directly with the client's stakeholders, validating ingestion, storage, and query latency on terabytes/day
- Enabled unified alerting across 7 networking-metric sources by building the Kubernetes-native ingestion layer for a Cisco observability PoC

Data Engineer

Sep 2022 - May 2023

NETMINDS

Remote

- Halved new-pipeline delivery time (4 days to 2) by rebuilding the platform around reusable ingestion templates; lifted loyalty engagement 10% with Spark / Scala pipelines feeding an AI-driven offers engine
- Shortened release cycles from 3 days to 1 by establishing Azure DevOps CI/CD for Databricks

Big Data Developer

Sep 2021 - Sep 2022

SOFTSERVE

Remote

- Built a Spark Structured Streaming system handling 50K+ events/second for real-time telecom network-map visualization and latency detection
- Translated ambiguous stakeholder requirements into service boundaries, SLAs, and data contracts; built demos used in design reviews that helped upsell the account

Data Engineer

Oct 2020 - Sep 2021

N-IX

Remote

- Secured sensitive datasets for a Fortune 500 company with a permissioned AWS data warehouse - Airflow, RBAC, Terraform CI/CD - plus a data-quality framework of 100+ automated checks

Education

M.Sc. & B.Sc., Artificial Intelligence - LPNU

Skills

Forward Deployed: On-Site Solution Design | Enterprise PoCs | Requirements to Data Contracts | Presales

AI / LLM: HuggingFace | LangChain | LLM-as-a-Judge Evaluation | Vector Retrieval | RAG | MLflow | Agentic Systems

Data & Streaming: SQL | Spark | Streaming | Databricks | Airflow | Kafka | Python | Scala | Flink | Snowflake

Cloud & DevOps: AWS (Sagemaker, Bedrock, Glue, EKS, S3) | Azure (ADF, Synapse) | MLOps | Kubernetes | Terraform